Autonomous Vehicles and the Gig Economy: Competition and Coexistence

Preliminary and incomplete, please do not circulate.*

Nicholas Buchholz† Jakub Kastl‡

September 3, 2026

Abstract

Autonomous driving is a major technological leap for taxi and ride-hail services, but threatens the livelihoods of millions of professional and part-time U.S. drivers. Using recent New York City data and the dynamic labor-supply model of Buchholz et al. (2026), in which drivers choose their hours endogenously, we evaluate two prospective regulatory programs for autonomous-vehicle (AV) deployment that do not harm incumbent drivers. The first is to tax AV revenue and rebate it to drivers, in the spirit of canonical proposals to tax automation and compensate displaced workers. We show it generally fails: AVs are close to pure substitutes for human drivers, so the tax recovers less than the earnings drivers lose. The second, which we propose, imposes no tax but restricts *where* AVs may operate, steering them to the trips that would otherwise leave drivers worst positioned for their next fare. Dispatching AVs to these trips keeps human driver earnings whole, raising ex-ante welfare even as AVs take a substantial share of trips. Regulating which trips AVs may serve can thus protect workers where taxing AV revenue alone cannot.

Keywords: autonomous vehicles, gig economy, ride-hailing, labor supply, dynamic discrete choice, AV regulation.

JEL Codes: J22, J24, L13, L92, R41.

*We are grateful to ... All errors are our own.

†Princeton University. Email: nbuchholz@princeton.edu.

‡Princeton University. Email: jkastl@princeton.edu.

1 Introduction

The rapid development of autonomous vehicle (AV) technology promises to transform urban transportation. Ride-hailing platforms such as Uber and Lyft have been explicit about their long-run intent to replace human drivers with autonomous fleets, and commercial robotaxi services have already launched in several major U.S. cities.[1] At the same time, AV and human driving are highly substitutable: AV entry and competition will reduce trip volume and earnings among taxi and ride-hail drivers. For the millions of workers who earn income as gig drivers, this prospect raises a central question for policy makers: is it realistic to permit the broad deployment of autonomous vehicles while also protecting workers from sudden job destruction? In this paper, we provide the first structurally grounded, empirically estimated answer for taxi and ride-hail labor markets.

Our primary concern is the major displacement risk along the AV transition path. In the long run, AV economics may support widespread deployment in lieu of human drivers. The policy question is how to get there. A transition that imposes sudden, concentrated earnings losses on the drivers now working invites purely protective responses: slowing deployment to spread the pain over time. We do not consider such a delay. Our interest is in complementary or self-funded mechanisms that let AV deployment proceed while leaving the incumbent driver cohort unharmed in both earnings and welfare, so that it exits only through ordinary attrition and retirement rather than displacement.

We consider two distinct regulatory programs. The first implements a public-finance solution, in the spirit of proposals to tax automation: unrestricted AV deployment is permitted, but AV revenues are taxed and the proceeds used to subsidize incumbent human driver labor. The second imposes no taxes but instead restricts where AVs may operate, steering deployment toward the trips human drivers value least. In neither program does the regulator control autonomous-vehicle penetration directly; it sets the levers (the AV tax and the subsidy it funds, or the permitted geographies), and the resulting penetration is an equilibrium consequence mediated by rider demand.

We begin with recent taxi and ride-hail market data from New York City. We do

[1] Waymo operates commercial driverless ride-hailing in Phoenix, San Francisco, Los Angeles, and Austin, and Tesla launched a robotaxi pilot in Austin in 2025; together these services have logged tens of millions of driverless miles. See Waymo (2024).

not have AV deployment or market performance data, which is at this time highly restricted to researchers. However, our question asks about outcomes for humans, for which we have both data and a robust, growing body of literature on driver opportunity costs and on the effects of regulation, competition, and congestion on earnings. We therefore treat AV deployment as the entry of competing robot taxi drivers, without any intrinsic car-specific welfare or earnings, and study their impact on human drivers.

AV regulation will affect many types of cities and markets, so our goal is to derive a set of principles that guide the feasibility and utility of these regulatory programs in environments beyond NYC, such as different geographies, market structures, and congestion levels. We also account for changing AV cost structures.

Our model of driver behavior adapts the dynamic labor supply model of Buchholz et al. (2026) (hereafter BSX). The model is attractive for our purposes because it allows for a simple and abstract representation of how demand, competition density, matching technology, and geography impact driver earnings through non-parametric equilibrium transition probabilities in earnings and hours. We analyze both the spatial-boundaries regime and tax-and-transfer regime by directly modeling their impact on these transition probabilities, accounting for network effects on the driver and passenger sides, geography and matching considerations, and endogenous congestion. Many of these effects are drawn from studies of the same market and time period: Frechette et al. (2019) for the demand response to waiting time and Mangrum and Molnar (2020) for vehicle-density congestion. We adopt the BSX model and estimate drivers' work-hours preferences directly on our data (Section 4.2). Counterfactual simulations then reduce to transforming the transition probability function to account for a variety of AV deployment and regulation scenarios and resolving for equilibrium earnings and hours worked, then repeating for different city structures, technology and congestion scenarios. The transition probability adjustments, representing the intra-daily likelihood of finding fares, serve as reduced-form representations of a wide range of counterfactuals, spanning alternative markets and geographic networks.

Our results contrast the two programs. The first, the canonical response to automation displacement, generally cannot generate enough revenue to replace drivers' lost income. A tax on autonomous-vehicle revenue recovers at most the automation's cost margin ($p - mc$), yet each autonomous trip displaces a human trip worth the full fare, so the rebate necessarily falls short of the driver's loss. Redistribution can

close this gap only if cheaper, faster service expands the market enough to stimulate demand, requiring demand to be highly responsive to waiting time. This is a narrow case, and we derive the conditions under which it holds. At autonomous-vehicle costs plausible today the condition fails: the tax-financed transfer cannot make incumbents whole.

The second program instead leverages the spatial spillovers to labor supply. Envisioning a unified platform that can dispatch AVs and human drivers, it allocates AVs to passenger trips whose destinations have the lowest forward-looking continuation values, the most unfavorable trips that leave human drivers facing the longest post-drop-off vacancy time. Holding demand fixed, human drivers then serve fewer but higher-quality trips. The net effect on earnings is not obvious ex ante: drivers face more competition on the routes that remain and higher congestion, while riders benefit from greater driver density and shorter waits, which in turn shifts demand out.

We find that this program can raise human driver earnings above the pre-AV baseline despite the entry of AV competition. In the 2025 ride-hail market, incumbent earnings peak at +2.5% and ex-ante welfare at +3.9% when autonomous vehicles are allowed about a tenth of the market's driving time, a fleet of roughly 500 cars serving one trip in twenty, net of the demand and congestion responses. Earnings remain at or above the baseline as the permitted fleet grows to serve a third of all trips.

The magnitude depends on market thickness: in thinner markets, where every trip already entails long vacancy time, the room for welfare-improving segmentation shrinks. Because the program's leverage comes from dispersion in continuation values rather than from anything particular to New York, the same model can be taken to other markets to read off how far AV deployment can go from a city's geography, dispatch technology, and vehicle density.

The two New York markets illustrate how the geography of that segmentation depends on the dispatch technology. In street hail, the program concentrates human service in the high-volume Manhattan core; in ride-hail it instead keeps human drivers across all of New York City while diverting autonomous vehicles to the most distant trips, such as those to New Jersey, Long Island, or farther.

A policy that can select *which* trips are automated protects workers where a tax on the surplus cannot. More broadly, this is a general strategy for the coexistence of human labor and automated capital in the presence of spatial spillovers: because

a worker's return on each task depends on where the last one left them, directing automation to the lowest-value tasks raises the value of the tasks that remain.

Related Literature

This paper contributes to three strands of the literature. The first studies market clearing and equilibrium in taxi and ride-hailing markets, on both the labor-supply and the spatial dimensions. On labor supply, Buchholz et al. (2026), whose framework we adopt, estimate a dynamic structural model of NYC medallion drivers' hours and wage elasticities, and Frechette et al. (2019) model driver entry as a dynamic game to study how search frictions and regulation shape wait times. On the spatial side, Buchholz (2022) characterizes the spatial search equilibrium of the NYC taxi market (where vacant taxis cluster given search frictions and pricing), motivating the reduced-form search-time channel through which AV entry lengthens human drivers' post-dropoff idle; we estimate that channel's elasticity from a shift-change regression on the 2012 data (Appendix A). Castillo (2025) and Rosaia (2025) study surge pricing and platform competition in related frameworks. We draw directly on two further studies: Buchholz et al. (2025) estimate consumer demand over price and wait time on the Liftago (Prague) platform, whose price-to-wait elasticity ratio disciplines the rider side of the tax-and-subsidy program (Section 5.1), and Mangrum and Molnar (2020) exploit the 2013 NYC Boro Taxi roll-out to estimate the marginal congestion of a taxi, which we use both as a regulatory analog and as the source for our vehicle-density congestion channel.

The second strand is the literature on the economics of autonomous vehicles and transportation market design. Ostrovsky and Schwarz (2018) is the closest antecedent. They study the interplay between autonomous transportation, carpooling, and road pricing, and characterize the conditions under which socially efficient outcomes can be implemented through appropriately designed tolls and cost-sharing rules.[2] Our paper is complementary: where Ostrovsky and Schwarz (2018) characterize the efficient AV equilibrium and the instruments that implement it, we ask what AV deployment is consistent with sustained incumbent driver earnings, taking as given the demand, supply, and dispatch primitives of an existing market. We hold

[2]Ostrovsky and Schwarz (2018) establish that a feasible, stable, budget-balanced, and market-clearing outcome exists and is socially efficient, blending matching-theoretic stability with Walrasian market-clearing tolls. Their analysis is theoretical and does not estimate demand or supply primitives from data.

the key instruments they emphasize, carpooling and optimal road pricing, fixed.

The third strand is the literature on automation and labor markets. Acemoglu and Restrepo (2020), Acemoglu and Restrepo (2018), and Autor et al. (2003) document the employment and wage effects of robotics and routine-task automation across manufacturing and services. Closest to our setting, Salman (2026) shows that commercial truck drivers respond to autonomous-vehicle testing even before displacement occurs, exiting licensing where AV exposure is high while remaining drivers extend their hours and earnings rise, reduced-form evidence that the driver labor-supply margin reacts to AV pressure. We take that margin as the object of a structural model: rather than measuring anticipatory responses to the threat of automation, we characterize the AV deployments consistent with sustained incumbent earnings under realized substitution. Our contribution brings this question into a platform setting where the structural labor-supply model of Buchholz et al. (2026) allows us to identify the slice of trips on which automation can be permitted without lowering incumbent driver earnings, and to characterize the dispatch-technology, density, and geography conditions on which the slice depends. Our second program speaks directly to the literature on taxing automation: whether the rents from labor-replacing technology can be taxed to compensate displaced workers (Korinek and Stiglitz, 2019; Costinot and Werning, 2023; Guerreiro et al., 2022; Thuemmel, 2023; Acemoglu et al., 2020); we develop that connection in Section 5.1. Chen et al. (2019) document the high intensive-margin elasticity of driver labor supply on Uber, consistent with the BSX continue/quit framework we use.

Roadmap. Section 2 describes the NYC datasets. Section 3 presents the BSX labor-supply model. Section 4 constructs and estimates the model for each market. Section 5.1 develops the tax-and-subsidy regulation (Program 1) and shows why redistribution generally cannot raise enough revenue to replace drivers' lost income. Section 5.2 then develops the spatial deployment regulation (Program 2), which formalizes AV entry as a substitution policy that transforms $\mathcal{F}$ and protects incumbent earnings by directing AV to the trips human drivers value least. Section 6 concludes.

2 Setting and Data

2.1 Data

Our supply-side primitives come from the New York City Taxi and Limousine Commission (TLC) trip records for yellow medallion taxis. Each record reports the time, location, distance, and fare of an individual trip, and a driver identifier that allows us to reconstruct complete driver shifts and the ordered sequence of trips within them. This shift-level reconstruction is what identifies the state-conditional trip-transition distribution $\mathcal{F}$ at the heart of the model: the distribution over the next trip's earnings, working time, and dropoff location, conditional on where a driver is and how far into a shift. We use the July–August 2012 window, before the September fare increase, and restrict attention to weekday (Monday–Thursday) shifts; the resulting sample contains roughly 646,000 trips across about 35,500 driver-shifts worked by about 15,500 distinct drivers. A representative weekday shift runs about 8.1 hours and earns roughly $245 in fares, an average of about $30 per hour.

To study the market that autonomous-vehicle entry will actually affect, we turn to the TLC's 2025 trip records for both yellow taxis and high-volume for-hire vehicles, the ride-hail services operated by Uber and Lyft. Ride-hail drivers are licensed by the TLC, and their trips are reported in the same record format, with request, pickup, and dropoff times and locations, trip distance, and driver pay. These records do not contain driver identifiers, however, so shifts cannot be reconstructed and $\mathcal{F}$ cannot be identified at the driver level; we therefore use trip-level distributions for 2025 and transfer the within-shift structure estimated from the 2012 panel (Section 4.1.2). Throughout, we aggregate the TLC's fine-grained taxi zones into a small number of geographic bins, defined where they are used.

The relative scale of the two markets has inverted since 2012. In 2025, ride-hail accounts for roughly 670,000 trips per day served by about 57,000 vehicles, against roughly 123,000 yellow trips per day on about 8,000 medallions in service on a typical day (a subset of the roughly 13,000 then licensed).[3] A representative ride-hail vehicle is active about 6.5 hours per day and a yellow taxi about 8.4 hours; average idle (non-trip) time is 14.9 minutes per ride-hail trip and 19.0 minutes per yellow trip,[4]

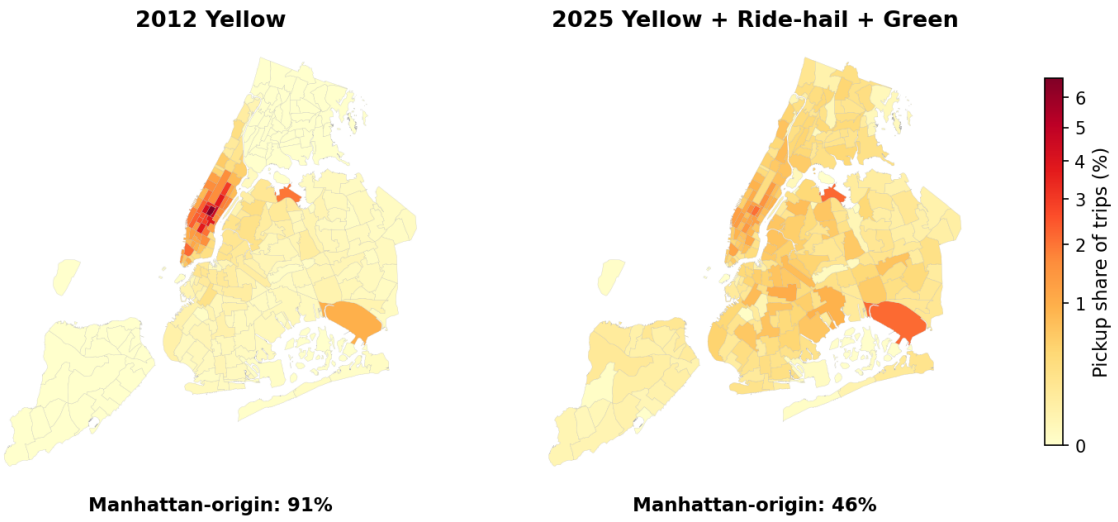
[3] All 2025 aggregates in this paragraph are from the NYC TLC industry indicators (NYC OpenData, dataset `v6kb-cqej`).

[4] Idle time per trip is not reported directly; we back it out from the same indicators as active

figures we use to anchor the matching calibration in Section 4.1.2.

The two markets also differ sharply in where trips originate (Figure 1). In 2012, street-hail pickups were overwhelmingly in Manhattan (91% of trips); by 2025 the combined yellow-and-ride-hail market is far more dispersed, with only 46% of pickups originating in Manhattan and the remainder spread across Brooklyn, Queens, and the Bronx. This broadening shifts where a driver-favorable segmentation can be drawn (Section 5.2).

Figure 1: Where trips originate, by taxi zone: 2012 versus 2025



Notes. Pickup share of trips at the TLC taxi-zone level: 2012 street-hail (left) versus 2025 yellow taxi plus high-volume for-hire (ride-hail) plus green (right; the 2025 panel combines all three modalities, trip-weighted). Each panel's shares sum to one; color uses a shared power-normalized scale. The Manhattan-origin share is 91% in 2012 and 46% in 2025.

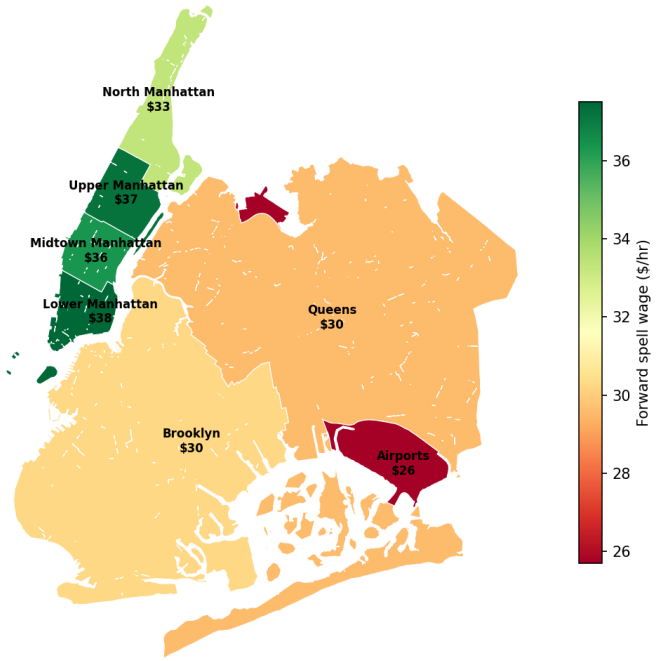
2.2 Heterogeneity in continuation value

The geographic segmentation fundamentally works because trips are widely dispersed in their continuation values, and the low tail is concentrated in predictable places. We measure this directly, without the structural model, by the *empirical forward spell*

minutes per trip less average trip duration. For ride-hail, $(6.5 \times 60)/(670{,}000/57{,}000) - 18.3 \approx 14.9$ minutes, where 670,000/57,000 is trips per vehicle per day and 18.3 minutes is average trip duration; for yellow, $(8.4 \times 60)/(123{,}000/8{,}200) - 14.7 \approx 19.0$ minutes. This is the total unpaid time between completed trips, which in a sequential single-rider market is the post-dropoff search to the next pickup.

wage: the average wage (in $/hr) a driver earns after a dropoff (and thus, vacant) in a given location until the completion of the next trip. A high forward spell wage means the dropoff leaves the driver well positioned for low search times and higher subsequent fare earnings; a low one means a long unpaid backhaul or idle search. Figure 2 maps this object across the seven-bin geography of the 2012 street-hail market.

Figure 2: Empirical forward spell wage by location, 2012 street-hail (medallion) market



Notes. The forward spell wage of a region is the average wage (in $/hr) a driver earns on the next spell after a dropoff there, computed model-free from the trip panel at the seven-bin geography used in the analysis; it is the data analog of the continuation value in the structural model (Section 3). Values fall from the Manhattan core (Lower Manhattan $38/hr, Midtown $36, Upper Manhattan $37) through the outer boroughs (Brooklyn and Queens $30) to the airports ($26), where a long unpaid return leaves drivers worst positioned; red denotes low values.

3 Model

To model labor supply among taxi and ride-hail drivers we adopt the dynamic stopping model of Buchholz et al. (2026) (BSX), and refer the reader to that paper for the model's full development, identification, and estimation. We restate here only the elements our counterfactuals require, emphasizing the transition function $\mathcal{F}$, the object our counterfactuals act on to represent the entry of autonomous-vehicle competition. The model is a single-agent optimal-stopping problem: taking the equilibrium earnings process as given, a driver chooses after each dropoff whether to continue searching or to quit for the day, so as to maximize expected shift utility; the value function below characterizes the driver's stopping rule.

A driver works a shift as a sequence of trips, deciding after each dropoff whether to continue searching or quit for the day. The state at a decision point is $s = (r, h, \ell)$: cumulative shift income r (in dollars), cumulative shift time h (in hours), and a discrete location ℓ. Flow utility from a state, with the coefficient on income normalized to one, is

$$U(r, h) \;=\; r \;+\; \theta_1\, h \;+\; \theta_2\, h^2, \tag{1}$$

with $\theta_1 > 0$ and $\theta_2 < 0$, so utility is hump-shaped in hours.

If the driver continues, he draws the next trip's increments to income and time and a new location from the transition function $\mathcal{F}$; if he quits, the shift ends and accumulated payoffs are collected. The continue/quit shocks are i.i.d. Type-I extreme value with scale σ, so the value function solves

$$V(s) \;=\; \sigma\, \log\Big(\exp\Big(U(s)/\sigma\Big) + \exp\Big(EV(s)/\sigma\Big)\Big) \;+\; \sigma\gamma, \tag{2}$$

where $EV(s) \;=\; \sum_{s'} \mathcal{F}(s' \mid s)\, V(s')$ is the continuation value, $\gamma \approx 0.5772$ is the Euler–Mascheroni constant, and there is no within-day discounting. Each driver type is summarized by three structural parameters, $(\theta_1, \theta_2, \sigma)$. A driver's realized shift length is an outcome of this stopping problem, determined by the continuation values and the realized continue/quit shocks, not by a static optimum of U.

The specification is parsimonious: three parameters per driver type, $(\theta_1, \theta_2, \sigma)$, with no hour-of-day or reference-dependence terms. BSX adopt this minimalism to show that a simple dynamic model reproduces behavior such as the negative end-of-shift wage elasticity, which static accounts ascribe to richer or reference-dependent

preferences; the parsimony is part of that argument.[5] We retain it here because it fits our data closely on untargeted moments (Section 4.2).

The absence of a clock-time term is less restrictive here than it may appear, because drivers are grouped by shift start time and start times are lumpy. Three-quarters of *am* shifts begin between 6:00 and 9:00 a.m., so within that group cumulative hours worked track the hour of day closely: conditional on h, the standard deviation of the clock hour is 1.2 hours, and the mean clock hour advances nearly one-for-one with h. For *am* drivers, who are 68% of shifts, h therefore proxies for time of day as well as for time on the job, and a separate hour-of-day term would be close to collinear with it. The *pm* group is more dispersed—its start times are bimodal, splitting between a midday and an early-evening cohort, and the same conditional standard deviation is 3.8 hours—which is also where the model's fit is weakest (Section 4.2).

3.1 The transition function $\mathcal{F}$

$\mathcal{F}$ is the central primitive of the model. Conditional on a driver's state $s = (r, h, \ell)$ it gives the joint distribution of the next trip's income increment Δr, time increment Δh (post-dropoff search plus time on the trip), and dropoff location ℓ'. Two features make it well suited to our counterfactuals.

First, $\mathcal{F}$ encodes the driver's labor-supply environment through its dependence on the state. Search and idle times grow with cumulative hours h (drivers tire and search less aggressively as the shift wears on), and continuation values fall after trips ending in low-density locations. Because $\mathcal{F}$ summarizes the equilibrium earnings process, itself the outcome of search and of the thick/thin externalities on both sides of the market, our counterfactuals can act on it directly, transforming it to simulate alternative market structures, dispatch technologies, and AV deployments without re-deriving equilibrium from primitives.

Second, $\mathcal{F}$ carries the heterogeneity the analysis requires. We index it by driver-shift type and by market conditions: drivers are split into *am*-shift starters versus the rest, and each shift is assigned to one of five daily-earnings quintiles capturing how

[5] This is the subject of their title: the negative end-of-shift wage elasticity, long read as evidence of daily income targeting, arises in their model from the dynamics of the stopping problem alone, with no reference point in the utility function. We carry that finding over rather than re-litigate it here.

productive driving is that day, the same exogenous earnings variation BSX exploit. $\mathcal{F}$ is thus a family $\mathcal{F}_{k,x}$, one transition function per (type, quintile) pair.

3.2 Welfare metrics

We report driver welfare with two measures. **Realized per-shift earnings** (our primary measure) is the dollar sum of trip fares averaged across simulated shifts; it is directly observable, and the object a regulator would target. The corroborating measure is V_{ante}, the Bellman value at the driver's realized initial state averaged across shifts, expressed in dollars; it is more sensitive to σ and to the curvature in h. The two need not peak at the same AV deployment: realized earnings stay at or above baseline across a range of AV shares while V_{ante} keeps rising with accumulated option value, so we treat realized earnings as the binding worker-protection constraint and read V_{ante} as corroborating evidence of an ex-ante welfare gain.

3.3 Entry

Assumption 1. The mass of potential drivers N is exogenous and does not respond to AV entry; only the intensive margin (when to end a shift) is endogenous.

We fix the extensive margin for two reasons. First, entry in NYC is strictly regulated: the medallion count is capped and for-hire-vehicle licensing has been actively restricted, so the pool of potential drivers is not free to adjust. Second, entry and exit tend to follow earnings, as drivers enter when driving pays well and leave when it does not (Hall et al., 2023), whereas our object of interest is precisely the upstream effect of AV entry on incumbent earnings, all else equal. Holding entry fixed isolates that effect and keeps the welfare quantification more interpretable across markets with different entry frictions. The fixed entrant pool can be regarded as the incumbent cohort: the drivers now working, whose earnings AV entry can disrupt suddenly and who would otherwise exit in the face of earnings losses. Endogenous entry would require additional assumptions about the pool of potential entrants and their outside options.

4 Empirical Analysis

This section takes the model to NYC data: for each market (the 2012 street-hail panel and the 2025 ride-hail records) we construct the equilibrium transition function $\mathcal{F}$, estimate the driver preferences, and fix the channels through which AV entry reshapes $\mathcal{F}$. Section 5 then evaluates the two regulatory programs against this estimated model.

4.1 Constructing the transition function $\mathcal{F}$

We discretize the state space and estimate a separate $\mathcal{F}_{k,x}$ for each driver-type/daily-quintile cell. Each $\mathcal{F}_{k,x}$ is a $K \times K$ transition matrix, where K is the number of discrete states: income bins $\times$ time bins $\times$ locations, equal to $22 \times 22 \times 7 = 3{,}388$ for the seven-bin 2012 geography and $22 \times 22 \times 10 = 4{,}840$ for the ten-bin 2025 geography (Section 4.1.2). Each matrix therefore has K^2 origin–destination state-pair entries, and we estimate one such matrix for each of the two driver types (*am* versus non-*am* start) and five daily-earnings quintiles. We restrict throughout to weekday (Mon–Thu) shifts and, relative to BSX, drop the owner/lease distinction, a margin of less interest in this study. Together these reductions cut computation roughly fourfold, allowing for a finer geography instead.

4.1.1 2012 Street-hail Baseline

The 2012 street-hail panel is our source for the labor-supply primitives, for two reasons. It carries driver identifiers, so we can reconstruct each driver's complete sequence of trips and estimate the state-conditional transition function directly. And it predates ride-hail, so the medallion market it describes is uncontaminated by for-hire-vehicle competition. In the 2025 data we lack access to driver identifiers even for yellow cabs (Section 4.1.2); this is why we recover street-hail driver preferences and the within-shift dynamics from the 2012 data. For the 2012 baseline we count transitions $(r, h, \ell) \to (r', h', \ell')$ directly from the observed trip sequences. We aggregate dropoff and pickup locations into seven NYC regions: Lower, Mid, Upper, and North Manhattan; Brooklyn; Queens; and the Airports (JFK and LaGuardia). Binning income and time into 22×22 cells over these regions (a $22 \times 22 \times 7$ state

space), $\mathcal{F}(s' \mid s)$ is the transition density.[6]

4.1.2 Constructing a 2025 Ride-hail Baseline

The 2025 public TLC data for for-hire vehicles (the ride-hail services of Uber and Lyft) carry no driver identifiers. We cannot reconstruct shifts, so direct estimation of a ride-hail $\mathcal{F}$ at the (r, h, ℓ) state level is infeasible. Instead, we split the two pieces of structure we need across the two datasets: the 2025 records provide the *spatial* primitives (location transitions and per-cell trip details) in their native geography, while the 2012 baseline provides the *temporal* structure of labor-supply, specifically how spell wages decline in cumulative work hours due to fatigue and break-taking.

We build this baseline at a ten-region geometry that adds to the seven 2012 regions (Section 4.1.1) three more regions where ride-hail does substantial business but the 2012 street-hail market did not: Bronx, Staten Island, and an Exterior region (Newark and outside-NYC).

We factor each transition as $\mathcal{F}(s, s') = P(\ell' \mid \ell)\, G(\Delta r, \Delta h \mid \ell, \ell', h)$. This imposes two assumptions: the next-location distribution $P(\ell' \mid \ell)$ depends only on the current location, not on accumulated earnings or hours; and the income and time increments G depend on the accumulated state only through cumulative hours h (the fatigue channel), not through cumulative income r. The spatial factor $P(\ell' \mid \ell)$ and the per-cell trip details (mean driver pay and baseline post-dropoff cycle time) are sampled from the 2025 ride-hail records; the temporal/labor-supply structure is then borrowed from the 2012 baseline, whose fatigue profile $\phi(h)$ scales each cell's baseline cycle time (See Appendix C for details). Trips that leave the zone grid receive a separate Exterior-return transition (an average 41.4-minute return-deadhead, with no fare and no fatigue scaling), because TLC-licensed ride-hail almost never picks up outside NYC. The fatigue profile $\phi(h)$ is common across the two modalities; the taste-for-hours parameter θ_1 is re-estimated to match the observed ride-hail shift length (Section 4.2), capturing the common use of ride-hail as a more flexible source of labor and earnings.

This common-fatigue assumption is the construction's strongest; testing it would

[6] In practice we estimate this via a Gaussian kernel: within each visited row a small weight (≈ 0.1) is shifted from the observed mass onto unobserved forward-adjacent states, with the kernel centered on the row's average income/time movement and decaying with grid distance. States never reached in the data are left empty, which caps shift length endogenously: a driver who lands in such a state must stop.

require the driver identifiers the public FHV data lack. The full details about how we transform $\mathcal{F}(s, s')$ (the kernel placement, the fatigue equation, and the Exterior special cases) are in Appendix C.

4.2 Estimation and model fit

We estimate the preference parameters $(\theta_1, \theta_2, \sigma)$ on the 2012 panel, following the methodology of BSX, so that the labor-supply model is disciplined by our sample and geography. Given the nonparametric transition function $\mathcal{F}$, a candidate θ implies a continuation probability $p_\theta(s)$ at every state through the Bellman equation (2); we choose θ to maximize the likelihood of the observed continue/quit decisions,

$$\hat{\theta} \;=\; \arg\max_{\theta} \; \sum_{i} \Big[q_i \log\Big(1 - p_\theta(s_i)\Big) + (1 - q_i) \log p_\theta(s_i) \Big] \;-\; \lambda\, \sigma^2, \tag{3}$$

where $q_i \in \{0, 1\}$ indicates a quit and the sum runs over decisions after the third hour of a shift (drivers essentially never stop before then). The small ridge penalty $\lambda\,\sigma^2$ regularizes the logit scale, which is otherwise weakly identified. We estimate separate parameters for the two shift types (*am*- versus *pm*-start), pooling owner-operator and lease drivers, on the weekday (Mon–Thu) sample. Estimation is a single maximum-likelihood problem per type and is described in Appendix A. For the 2025 ride-hail market the public data carry no driver identifiers, so preferences cannot be estimated by this likelihood. Directly using the 2012 street-hail preference estimates over-predicts ride-hail shift length, since ride-hail drivers are largely part-time and work shorter shifts. We instead estimate the ride-hail taste for hours from the one aggregate moment the public data support: a single downward shift in θ_1 chosen so that the model reproduces the average ride-hail shift length (Appendix A). The resulting shift is sizable, consistent with a higher outside option for a multi-homing, part-time population.

Because the estimator targets only the continue/quit decisions, the resulting shift length, earnings, and trip counts are untargeted: so their alignment with the data serves as a direct test of fit. Table 1 reports the comparison. The model reproduces the average shift length (8.0 vs 8.1 hours) and hourly earnings rate ($29 vs $30) closely, and the by-type split (longer, higher-earning *am* shifts) in the right direction. As a further, out-of-sample check, holding θ fixed and feeding the model the

Table 1: Model fit: untargeted moments, 2012 baseline. The estimator matches only the continue/quit decisions; shift length, earnings, and trip counts are outcomes. Rates are ratios of means.

	am shift		**pm shift**		**All**	
Moment	Data	Model	Data	Model	Data	Model
Shift length (hours)	8.48	8.60	7.26	6.64	8.09	8.00
Earnings per shift ($)	252	245	230	206	245	233
Earnings per active hour ($)	29.7	28.5	31.7	31.0	30.3	29.1
Trips per shift	20.3	22.1	13.7	16.7	18.2	20.5

directly estimated post-period $\mathcal{F}$ reproduces the change in driver behavior across the September 2012 fare increase (Appendix B), direct evidence for the transportability the counterfactuals rely on.

5 Program Evaluation

We evaluate two regulatory programs against the estimated model. We begin with Program 1 (tax-and-subsidy), the canonical response to automation displacement (tax the technology's revenue and redistribute it to affected workers), and show that it generally cannot raise enough revenue to offset drivers' lost income at current AV costs. We then turn to Program 2 (spatial restriction), which directs autonomous vehicles to the trips human drivers value least; it keeps incumbent earnings whole and raises ex-ante welfare where the tax cannot.

5.1 AV Program 1: Tax-and-Subsidy

We first consider taxing AV trips, without any spatial restrictions, and rebating the tax back to taxi and ride-hail drivers in proportion to their trip earnings. This policy reflects a popular policy idea to shift tax burdens from labor towards capital in the age of artificial intelligence (Korinek and Stiglitz, 2019; Costinot and Werning, 2023; Guerreiro et al., 2022), e.g., a tax on labor-saving AI systems to subsidize workers facing reduced work hours. The taxi-driver environment is a particularly transparent version of such a policy, since unlike many workplace settings, AV taxi rides are essentially pure substitutes for human taxi rides in the service provided. In addition, general equilibrium issues related to two-sided market externalities and congestion

are all well-defined.

The specific policy we consider is this: the regulator taxes a share τ of all AV trip revenues and transfers it to incumbent drivers. As a baseline, we continue to assume that service prices are identical and dispatched by the same platform as the FHVs, offering a constant 5-minute waiting time. Nevertheless, there is an AV participation constraint that the after-tax fare is higher than the average cost, capping the tax at $\tau \leq \tau_{\max} = 1 - c_{AV}/p$.

We first show that the above constraint will almost always bind given today's technology, rendering the tax-and-subsidy policy generally infeasible until AV costs fall low enough. It fails to bind only when demand is highly elastic to waiting time and average costs are low.

5.1.1 A simple theory of tax-financed compensation

With a fixed pool of incumbent drivers, let R denote their baseline fare revenue and W their total income after AV entry: the fares they still serve plus the per-trip subsidy they receive. Relative to baseline, W satisfies the accounting identity

$$\frac{W}{R} \;=\; d\Big[\,(1-\sigma) + \tau\rho\,\sigma\,\Big] \;=\; d\Big[\,1 - \sigma(1-\tau\rho)\,\Big],$$

where σ is the AV share of trips, d is total trips relative to baseline (demand expansion), $\rho = p^{AV}/p^{H}$, and τ is the share of AV gross revenue taxed and transferred in full to drivers. AV participation requires the operator to cover its cost, bounding the tax at $\tau \leq \tau_{\max} = 1 - c_{AV}/p$. Write $G \equiv W/R - 1$. The baseline treats the autonomous and human services as *perfect substitutes* (identical at a common fare, with infinite cross-price elasticity), so any AV price premium drives its share to zero (Proposition 2); the structural rider-choice model (Section 5.1.2) relaxes this to imperfect substitution.

Proposition 1 (inelastic demand, fixed prices). With $d = 1$ and $\rho = 1$, $G = -\sigma(1-\tau) < 0$ for every feasible $\tau \leq \tau_{\max} < 1$. Each displaced trip costs drivers the full fare but returns only the fraction τ of it; at the ceiling the uncompensable loss per trip is exactly c_{AV}, the automation's resource cost. $G = 0$ would require $\tau = 1$ (taxing the entire fare), which no operating AV permits.

Proposition 2 (flexible prices, wait-inelastic demand). Raising ρ so that the tax per AV trip exceeds the lost fare ($\tau\rho \geq 1$) requires $\rho \geq 1/\tau > 1$. But with no quality advantage an AV priced above the human fare is strictly dominated, so $\sigma(\rho) = 0$ for $\rho > 1$: the tax base vanishes before it can compensate. The only configurations are $\rho = 1$ (Proposition 1, negative) and $\rho > 1$ (no AV). There is no driver-positive interior.

Proposition 3 (the bound under wait-elastic demand). AV sustains a role only by offering lower waits, which expands demand. For a marginal deployment s under fare equalization ($\rho = 1$) at the maximal tax, capacity rises by $s(1-\phi)$ where ϕ is the human-hours contraction, $d \approx 1 + \eta|\varepsilon_w|(1-\phi)\,s$, and $\sigma \approx s$, so

$$G \;\approx\; s\Big[\,\eta|\varepsilon_w|(1-\phi) - c_{AV}/p\,\Big].$$

With no hours contraction the compensation scheme breaks even iff

$$\boxed{\eta\,|\varepsilon_w| \;\geq\; c_{AV}/p} \tag{4}$$

The composite wait-quality demand elasticity must exceed the automation's cost share. At our NYC 2011–12 calibration $\eta|\varepsilon_w| = 0.265 \times 1.225 = 0.325$, against $c_{AV}/p \in [0.34, 0.47]$ (marginal to full cost): the condition fails, and the market-clearing limit ($\phi \to 1$) drives $G \to -s\,c_{AV}/p$, negative at every scale, exactly the structural result in Section 5.1.2.

Lower waiting time is a service-quality improvement, so the autonomous vehicle is not purely cost-reducing automation but quality-improving automation that can enlarge the market. Compensation becomes feasible when that quality gain raises aggregate demand faster than automation widens the per-driver earnings gap. Cost-reducing automation, which only displaces the human trip, can never fund the transfer this way; sufficiently quality-improving automation can.

Signing the omitted effects. The marginal identity lets us sign each channel we do not model (Table 2). Every effect that moves G down (congestion, human-hours contraction) only widens the gap; the only levers that close it raise $\eta|\varepsilon_w|$ (a higher-throughput matching technology or more wait-elastic riders, i.e. a different market) or lower c_{AV}/p (cheaper AVs over time). This yields the paper's central policy statement:

Table 2: Direction of each unmodeled channel on the marginal driver gain G', from the accounting identity.

Omitted effect	Sign on G	Mechanism
Composite elasticity $\eta\|\varepsilon_w\|$	$+$	only lever that helps; needs a different market
AV cost share c_{AV}/p falling	$+$	raises $\tau_{\max}$ and lowers the threshold (dynamic escape)
Human-hours contraction ϕ	$-$	thins the market; the market-clearing limit
Congestion β_{cong}	$-$	AV density slows trips, shrinks demand expansion
Selection on diversion	0	neutral under random allocation
Tax rate τ	$+$	helps, but hard-capped at $\tau_{\max} = 1 - c_{AV}/p$

rent-financed compensation of displaced workers is infeasible at current automation costs and becomes feasible only as the technology's cost share falls, just as the rent grows large enough to fund it. The corollary for robot-tax proposals is stark: the tax suffices only when job creation $(d-1)$ is high, which is exactly the state of the world in which compensation is least needed.

A price floor without redistribution ($\rho > 1$, $\tau = 0$) cannot protect incumbents (Proposition 1), and a premium high enough for the tax to exceed the displaced fare cannot be sustained by rider choice (Proposition 2). The instrument we therefore evaluate is equal pricing ($\rho = 1$, the autonomous vehicle charges the human fare) with the cost margin taxed at the ceiling and recycled. Penetration is not chosen directly by the regulator: under a price floor it emerges from rider choice given (ρ, w^{AV}), and under equal pricing it is set by the size of the deployed fleet.

5.1.2 Quantifying the tax-and-subsidy equilibrium

To take the policy to the data we relax the perfect-substitutes baseline of Section 5.1.1. Riders on each origin–destination cell choose between the human and AV service in a binary logit over fares and waits, with value-of-time heterogeneity, making the two services *imperfect* substitutes with a finite cross-price elasticity. The preference inputs are elasticities estimated in the same market and period as the labor-supply panel. From Buchholz (2022) we take distance-graded price elasticities, $\varepsilon_p = -2.22$ for trips under two miles rising to $\varepsilon_p = -1.07$ beyond six miles, identified off the September 2012 tariff change. From Frechette et al. (2019) we take the wait-time elasticity, $\varepsilon_w = -1.225$.[7] And from Buchholz et al. (2025) we take a value-of-time distribution,

[7] Throughout we write ε_w for the wait-time demand elasticity (estimated by Frechette et al., 2019) and ε_p for the price elasticity (Buchholz, 2022); the composite $\eta\,|\varepsilon_w|$ in the break-even condition (4)

pinned by the scale-free ratio of wait to price elasticities (0.10 to 0.25), a median of roughly \$14.6/hour at NYC fares. The AV share σ_c, human waits, and trip volumes are solved jointly as a fixed point: AV diversion thins the human market, search times stretch through the same matching elasticity η as in Program 2, and demand responds to the resulting menu, so welfare numbers are directly comparable across the two programs. Tax revenue is rebated to incumbent drivers. We report a per-license lump-sum rebate, which induces no behavioral distortion; a per-fare piece rate (roughly +35% per trip at τ near its ceiling) transfers the same total under market clearing and yields the same driver income, so the choice of rebate design does not affect the conclusion.

We solve the equilibrium under fare equalization (the autonomous vehicle charges the prevailing human fare) with the tax set at its ceiling, so the entire cost margin is transferred to drivers. We impose market clearing directly: with a fixed pool of N drivers, incumbents can serve only the $(1-\sigma)d$ trips autonomous vehicles leave them, so per-driver fare income is pinned at that clearing volume rather than inferred from the search elasticity extrapolated far out of its estimation range. Figure 3 reports the result. Driver income net of the rebate is negative at every fleet size and falls monotonically with autonomous-vehicle penetration; the optimal tax-financed fleet is zero, confirming the analytical break-even condition (4). At the relevant small deployments the added AV capacity barely expands total trips ($d \approx 1$), so the autonomous vehicle is close to a pure substitute. A tax then returns at most the fraction τ of each displaced fare while drivers lose the whole fare, and net income falls by approximately $\sigma\,(c_{AV}/p)$.

This is the corner solution of a more general balancing rule, shown in the right panel. When the autonomous vehicle adds net capacity rather than merely replacing human hours, total trips expand and the driver gain $G(\sigma) = d(\sigma)\,(1 - \sigma\, c_{AV}/p) - 1$ is a concave, saturating benefit multiplied by a linearly declining displacement cost. It has an interior maximum,

$$\sigma^* = \frac{\eta|\varepsilon_w| - c_{AV}/p}{(c_{AV}/p)\,(1+\eta|\varepsilon_w|)},$$

positive only when the break-even condition (4) holds ($\eta|\varepsilon_w| > c_{AV}/p$). At the present calibration the numerator is negative and the optimum is at zero. As the autonomous-

pairs ε_w with the search-time elasticity η.

vehicle cost falls below the composite elasticity $\eta|\varepsilon_w| \approx 0.33$, an interior optimal fleet appears and grows: a cost share of 0.30 implies $\sigma^* \approx 0.06$, and 0.20 implies $\sigma^* \approx 0.47$. An interior optimum thus requires both a low enough cost and a demand expansion that survives the contraction of human hours, the latter governed by the parameter identified from the September 2012 fare change. The threshold cost does not depend on driver preferences: since η, ε_w, and the fare pin down $c^* = \eta|\varepsilon_w|\, p$, a revenue-neutral tax breaks even exactly when the autonomous-vehicle cost share falls to $c^*/p = \eta|\varepsilon_w| = 0.325$, that is, to 32.5% of the fare. Today's cost share of 0.34 to 0.47 (marginal to full cost) sits above this threshold, so the optimal tax-financed fleet is zero; the gap is decisive at full cost but a knife-edge at the marginal-cost end (0.325 against 0.34). In annual terms, an autonomous vehicle costing \$72,560 at full cost would have to reach about \$50,400, a reduction of roughly 30%, before the tax could hold incumbents harmless.

The same identity shows that an untaxed price floor cannot substitute for the tax. With $\tau = 0$ the floor alone has no interior optimum: incumbent earnings fall by 16–17% across the entire band $\rho \in [0.5, 1.5]$ and approach the no-AV baseline only as ρ grows large enough to drive $\sigma \to 0$. Demand at NYC fare levels is price-inelastic enough that the share responds weakly to the floor (σ falls only from 0.76 to 0.41 as ρ rises from 0.5 to 1.5), and rider choice allocates AV trips by fare, which is unrelated to the post-dropoff continuation values that make geographic segmentation welfare-improving. Relative to that segmentation, Program 2 raises incumbent earnings by 3.2% at a penetration of 4.7% with no transfer at all, because it selects which trips are displaced, whereas the tax-and-subsidy instrument, even with the entire cost margin recycled, cannot make incumbents whole at the autonomous-vehicle costs observed today.

5.1.3 Relation to the literature on taxing automation

Our finding speaks to an active debate over whether automation (and now AI more broadly) should be taxed to protect the workers it displaces. The comparison throughout runs through the break-even condition (4), $\eta\,|\varepsilon_w| \geq c_{AV}/p$, derived in Section 5.1.1.

The size of the tax. The optimal-tax literature concludes that the corrective tax on automation should be small and temporary. Costinot and Werning (2023) find

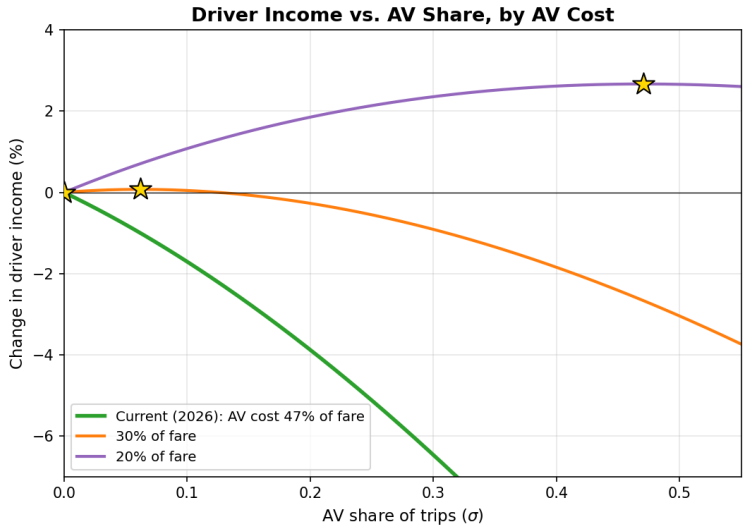


Figure 3: Tax-and-subsidy: change in incumbent driver income as a function of the autonomous-vehicle share of trips, for a ladder of autonomous-vehicle costs. The autonomous vehicle charges the prevailing human fare, and the full cost margin is taxed and returned to drivers as a per-trip rebate; each curve fixes an AV cost, and stars mark the optimal share. At current (2026) cost the optimal share is zero (green); at lower costs an interior optimum appears and grows. All curves assume net capacity expansion, so only AV cost varies across them.

an optimal robot tax of one to four percent that declines as automation deepens; Guerreiro et al. (2022) find a tax that is positive only during the transition and falls to zero once the displaced cohort leaves the workforce; Thuemmel (2023) finds the optimal tax is negative (a subsidy) while automation is expensive, turning slightly positive only as it cheapens. Our result is the opposite: merely holding the displaced drivers harmless requires taxing the entire cost margin (the largest feasible tax), and even that fails at current AV costs. The difference is in the question asked: the literature asks what tax is socially optimal, whereas we ask what tax would *fully compensate one displaced group*, a far more demanding standard that a small tax cannot meet. We agree on the transitory logic: the drivers to be protected are an existing cohort, and once they retire there is no one to replace and no one to compensate, so the warranted transfer (and the tax that funds it) shrinks over time, as in Guerreiro et al. (2022) and Thuemmel (2023).

Productivity, cost, and congestion. In the aggregate models a larger productivity gain weakens the case for a protective tax: it strengthens the productivity effect

of Acemoglu and Restrepo (2019), raising labor demand elsewhere. In our setting the productivity gain is concrete (a lower autonomous-vehicle cost relative to the fare), and it is what relaxes the break-even condition, so a larger gain makes compensation feasible, not merely less necessary; as c_{AV}/p falls the condition is eventually met. But the gain is not unambiguously benign here. More autonomous vehicles on the road add congestion (Mangrum and Molnar, 2020), which erodes the demand expansion that compensation relies on, unless the technology also reduces congestion, through coordination or higher occupancy, a possibility on which our model takes no stand.

When rent taxation can compensate. Korinek and Stiglitz (2019) show that technical change can be made Pareto-improving when its gains accrue as rents that can be taxed without distorting behavior; our analysis asks when that condition holds in a specific market. The taxable rent per displaced trip is the fare net of the autonomous-vehicle cost, $p - c_{AV}$, while the driver's loss is the full fare p, so rent taxation alone restores drivers only in the limit $c_{AV} \to 0$ (Proposition 1); a strict gain requires the productivity gain to expand the market, the productivity and reinstatement effects of Acemoglu and Restrepo (2018, 2020). The cost advantage can fund the transfer or expand the market, but not both: a high AV fare preserves a taxable rent yet blocks expansion, while a low fare expands the market yet leaves no rent to tax. Compensation through taxation therefore succeeds only where the productivity gain is large enough to expand the market substantially, which is also where displacement is most easily offset with no transfer at all.

5.2 AV Program 2: Spatial Restrictions

Under this program the regulator restricts *where* autonomous vehicles may operate, dispatching them to the trips that are least valuable for human drivers to keep. We implement this as a policy that selects which of the $K \times K$ origin-destination pair cells of the transition function $\mathcal{F}$ the autonomous fleet serves: the served cells are removed from the human drivers' $\mathcal{F}$, and the policy question is how many origin-destination pairs can be carved out before incumbent earnings fall below the no-AV baseline. Describing the implementation requires two choices: (1) which trips the AV serves (the origin-destination ranking rule, Section 5.2.1) and (2) how much it serves (a trip-hour capacity B, Section 5.2.3). Given these, we recompute how the diverted trips reshape the remaining drivers' search (the calibrated channels in Table 3).

Two assumptions frame the exercise. First, we do not model competing platforms: we envision a single platform that dispatches both human-driven and autonomous vehicles and chooses which car type serves each trip, so the spatial restriction operates through that one dispatcher rather than through competition between an autonomous-vehicle firm and a human-driver firm. Second, we assume an autonomous vehicle is always available within five minutes ($w_{AV} = 5$ min), reflecting a fleet dense enough that AV wait time is never the binding margin; this wait enters the rider-side calculation in the demand-response channel and returns as a regulated lever in Program 1.

5.2.1 Classifying locations to be served by AV

In both baselines we rank origin–destination cells by a single location-level score V_d for the destination d: the forward-looking dollars-per-hour value of a dropoff there, with the spatial restriction directing AV to the lowest-valued cells first. We score at the location level because that is the feasible policy lever: a regulator restricts where AV operates and cannot condition on an individual driver's accumulated hours and income, which differ across drivers and over a shift.

The two baselines differ only in how V_d is computed, because they afford different data. In the 2012 street-hail baseline we observe the full state-conditional $\mathcal{F}$, so V_d is the mass-weighted average, over the dropoff states $s' = (r, h, d)$ at location d, of the per-dropoff spell wage $\bar{w}(s')$, the expected per-hour earnings of the next trip:

$$V_d \;=\; \frac{\sum_{s':\,\ell(s')=d} m(s')\,\bar{w}(s')}{\sum_{s':\,\ell(s')=d} m(s')}, \qquad \bar{w}(s') \;\equiv\; \mathbb{E}_{\mathcal{F}}\left[\frac{\Delta r}{\Delta h}\,\middle|\, s'\right] \tag{5}$$

(with Δh in hours and $m(s')$ the empirical mass of dropoffs at s'). Forming V_d this way lets the state-conditional $\mathcal{F}$ absorb how forward value varies with accumulated hours and income before we integrate them out. In the 2025 ride-hail baseline no state-conditional $\mathcal{F}$ is available (the public data lack driver identifiers) so the same location score is formed directly from the trip records as a volume-weighted rate:

$$V_d \;=\; \frac{\sum_{d'} n_{d,d'}\,\overline{\text{earn}}_{d,d'}}{\sum_{d'} n_{d,d'}\left(\overline{\text{trip}}_{d,d'} + \text{idle}_{d'}\right)}, \tag{6}$$

the mean dollars per hour of trips originating at d (trip time and idle in hours). Both

are dollars per hour at d and only their ordering enters the deployment; the 2025 score is the coarsening of the 2012 one that the data permit.

5.2.2 Equilibrium Calibration

Our counterfactuals operate entirely by transforming $\mathcal{F}$. Table 3 collects the channels through which AV entry reshapes $\mathcal{F}$, the calibrated value each carries, and where its derivation is given. The AV deployment itself removes the served cells from $\mathcal{F}$; three further channels modify the remaining drivers' per-trip economics; and the AV cost converts a deployment into a fleet count and a dollar figure.

Table 3: Equilibrium calibration: the channels through which AV entry reshapes the transition function $\mathcal{F}$, with their calibrated values.

Adjustment	Effect on $\mathcal{F}$	Calibration (value, source)	Detail
AV substitution	Removes the AV-served cells and renormalizes the remaining row mass	Capacity B (policy instrument)	§5.2.3
Search-time stretch	Lengthens Δh as humans lose rider density	$\eta = 0.265$ (authors' calculation)	App. A
Demand response	Recovers part of the lost demand via shorter rider waits	$\varepsilon_w = -1.225$ (Frechette et al., 2019)	App. D
Congestion	Lengthens Δh with added vehicle density	$\beta_{\text{cong}} = 0.358$ (Mangrum and Molnar, 2020)	App. D
AV cost	Cost side; maps capacity B to a fleet count and dollar cost	\$72,560/AV/yr (Bösch et al., 2018)	App. A

5.2.3 Mechanism

Autonomous vehicles serve the origin–destination cells of $\mathcal{F}$ in ascending order of their continuation-value score (the ranking rules above), so the cells worst for human drivers are handed to AV first. Serving a cell removes it from the human drivers' $\mathcal{F}$; the remaining row mass is renormalized and the Bellman is re-solved, so drivers re-optimize their stopping rule against the trips that are left.

We trace outcomes as the size of the AV deployment grows from zero, re-solving the Bellman on the new $\mathcal{F}$ at each capacity on a grid $B/B_{\max} \in \{0, 0.025, \ldots, 0.7\}$. Two quantities describe that size in intuitive terms: the share of trips served by AV, σ_{AV}, and the implied fleet, the number of autonomous vehicles required. Because an autonomous vehicle's cost scales with the hours it drives rather than the number of trips it makes, the deployment is constrained by its total trip-hours: AV absorbs cells up to a given trip-hour capacity, which we write as $B/B_{\max}$ (its share of the whole market's daily trip-hours, $B_{\max}$), and which maps directly to a fleet of about $B/6.6$ vehicles (a 12-hour duty cycle; see Appendix A). The three measures of trip share, fleet, and capacity, rise together. Along with trip share and fleet we report $B/B_{\max}$, the share of total trip-hours performed by autonomous vehicles, which rises as automated cars are added to the road.

The *regulated AV slice* is the largest deployment at which expected per-shift human earnings stay at or above the no-AV baseline.

5.2.4 Program 2 Results

First we examine AV deployment in the geographically concentrated 2012 market. Across a wide band of deployments (from $B/B_{\max} = 0.05$ out to about 0.47), realized per-shift earnings stay at or above the no-AV baseline, so incumbents are not harmed. The earnings gain itself is small, peaking at about +0.7% near $B/B_{\max} = 0.35$ ($\sigma_{AV} = 0.21$); ex-ante welfare rises more and keeps rising with the deployment, reaching +7.0% at $B/B_{\max} = 0.45$ (a fleet of about 1,636 autonomous vehicles, $\sigma_{AV} = 0.32$, a third of trips). The result is coexistence with a modest gain: automation removes the "bad" destinations, and drivers, better positioned after each trip, are weakly better off.

Next we run the same simulator on the expanded 2025 ride-hail market, with the taste-for-hours parameter re-estimated to match the observed ride-hail shift length (Section 4.2). The resulting baseline reproduces about \$174 per shift, 6.2 hours, and \$28 per active hour, against TLC indicator values near \$200, 6.5 hours, and \$30.

At the default search-time elasticity $\eta = 0.265$, incumbent earnings peak at $B/B_{\max} = 0.10$: a fleet of about 509 autonomous vehicles serving 4.7% of trips, with $\Delta\text{earn} = +2.54\%$ and $\Delta V_{\text{ante}} = +3.87\%$, already net of the demand-response and congestion channels (Appendix E); Figure 5 traces the gain as the AV share rises. The earnings gain is larger here than in 2012 because ride-hail drivers work

Table 4: Labor-optimized AV deployment in the two markets.

	2012 baseline	**2025 ride-hail**
Earnings-max policy:		
$B/B_{\max}$	0.35	0.10
σ_{AV}	0.21	0.05
AV fleet size	1,270	509
Δ earnings	+0.74%	+2.54%
ΔV_{ante}	+5.02%	+3.87%
Earnings-positive band:		
$B/B_{\max}$ range	$[0.05, 0.47]$	$[0, 0.38]$
Max AV fleet	1,700	1,900
Max σ_{AV}	0.33	0.34

shorter baseline shifts, leaving more room to extend. Earnings stay at or above baseline from $B/B_{\max} = 0$ out to about 0.375, up to roughly 1,900 vehicles and 34% of trips. Beyond this level, automation begins absorbing high-value within-city trips. At peak, the automated trips are predominantly out-of-city flows: roughly 60% of the automated trip-hours, together with the lowest-value within-city trips.

In both markets the gain runs through the labor-supply margin rather than the wage rate: the per-active-hour rate is essentially unchanged (about $28 at the 2025 peak, before and after). With automation taking the lowest-ranked trips, the continuation value rises at every state a driver might reach, so drivers voluntarily extend their shifts and earn more by working longer (Δhours/hours $\approx$ Δearn/earn).

The reported gain is for the fixed simulated pool of incumbents (Assumption 1). Under free entry (the Hall et al. (2023) logic), positive incumbent earnings would draw new drivers and dissipate the rent over time. Either way the labor-optimization criterion is met at the peak: incumbents are no worse off than the no-AV baseline. The +3% is thus an upper bound on the per-incumbent gain, with the split between incumbents and entrants depending on an entry margin we do not model.

Which trips end up automated. Both markets sort by destination, but where the line falls depends on dispatch technology (Figure 4). In 2012, street-hail cruising makes non-central destinations much less productive, so human drivers keep the Midtown–Lower-Manhattan core and automation takes the rest of the city (30 of the 49 origin–destination cells go to automation, 18 stay with drivers). In 2025, fast

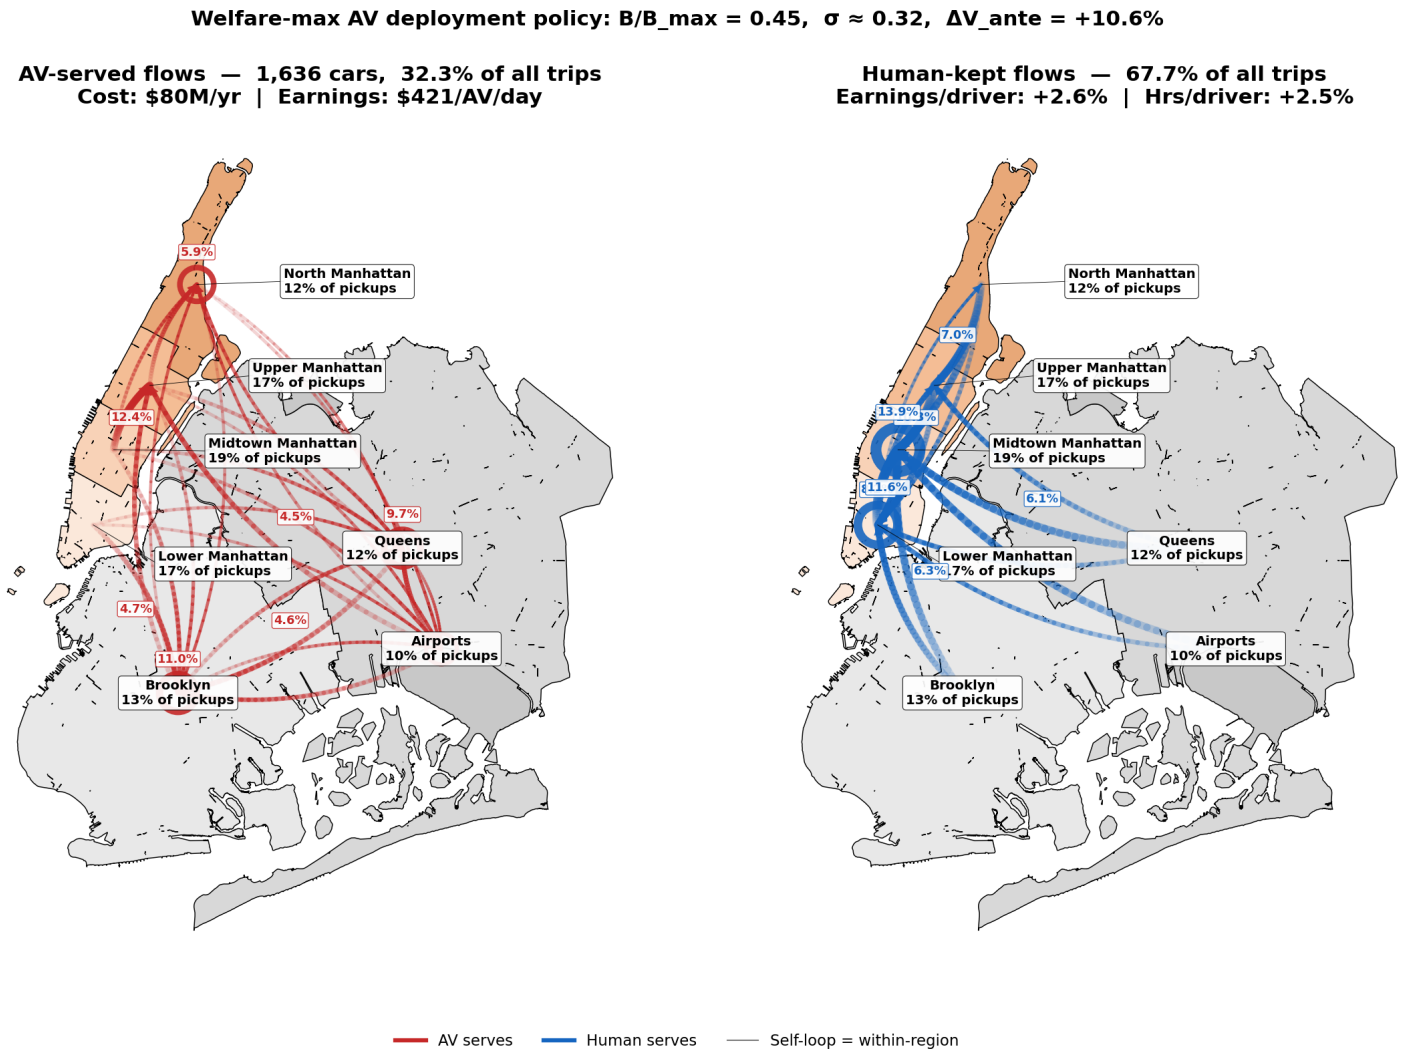


Figure 4: Earnings-max AV deployment in the two markets: automation (red) takes the lowest-value trips, human drivers (blue) keep the rest. **(a)** In 2012 street-hail ($B/B_{\max} = 0.45$), automation serves every trip dropping outside the Midtown / Lower-Manhattan core and human drivers keep the core.

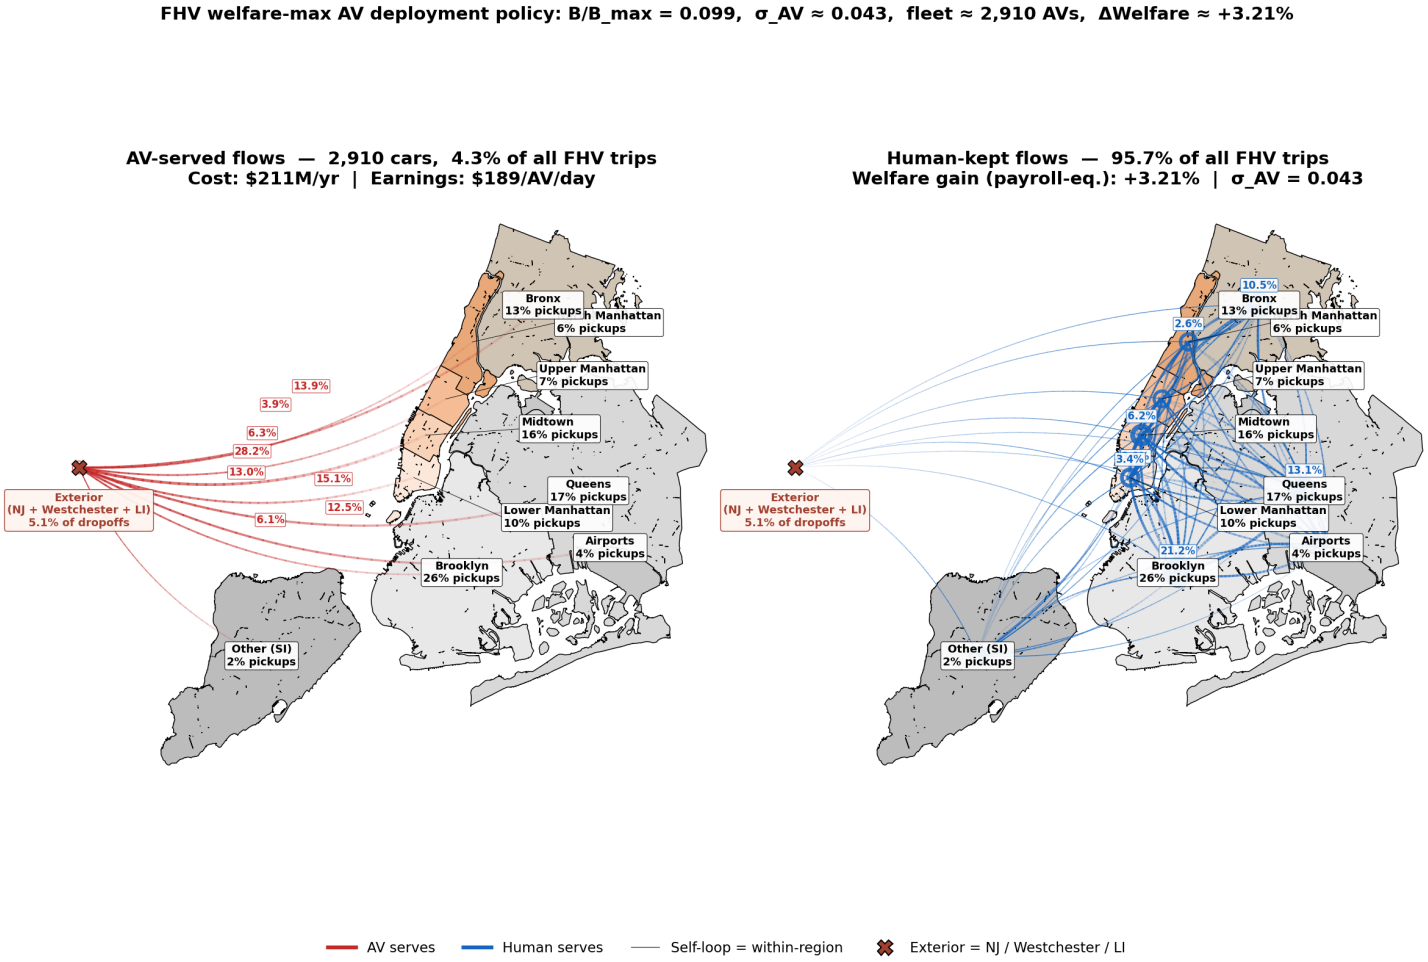


Figure 4: Earnings-max AV deployment in the two markets (continued). **(b)** In 2025 ride-hail (static surrogate, $B/B_{\max} = 0.10$, $\sigma_{AV} = 0.05$), every automated trip terminates at the Exterior (off-canvas, west) and human drivers keep every within-city trip.

Table 5: FHV destination continuation values V_d.

Destination d	V_d **(\$/hr)**	**Below \$35.4 mean?**
Airports	54.30	×
Midtown	39.19	×
Lower Manhattan	37.92	×
Upper Manhattan	37.49	×
Other (Staten Island)	34.23	✓
North Manhattan	33.38	✓
Brooklyn	31.87	✓
Queens	31.64	✓
Bronx	30.21	✓
Exterior	23.57	✓

matching has compressed forward values within the city, so drivers keep all of New York and automation is pushed to trips that leave it, to Newark and beyond the five boroughs. Those are the least attractive trips to keep: they end far from the next fare, and the return is on average an unpaid 40-minute backhaul with little chance of a pickup en route (some 1.4 million such trips end outside the city in our three-month sample, against essentially none that begin there). Table 5 ranks destinations by forward value; out-of-city trips sit at the bottom, near \$24/hr against a \$35 citywide average.

The same logic indicates where the policy helps most: where demand is concentrated in a dense core with lower-value outskirts (the 2012 pattern, and the 2025 pattern at the city's edge), and least where demand is spread uniformly, leaving no clearly worse trips to hand off. Geographic incidence and the size of the permitted fleet thus turn on both dispatch technology and the shape of urban demand.

5.2.5 Robustness

We explore the earnings results sensitivity along four dimensions, with details in Appendix E (market density and the demand and congestion channels) and Appendix A.1 (the search-time elasticity η). At observed market density the earnings-max policy stays positive throughout: adding the demand-response and congestion channels leaves the peak gain essentially unchanged (the two roughly cancel), and varying η across its plausible range moves the gain between about +2.2% and +3.1%. Market density is the lever that bites: the peak gain falls from +2.5% at full density

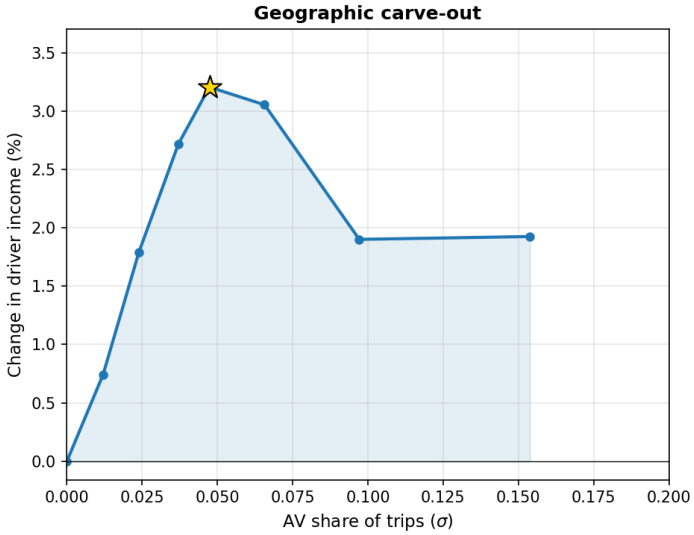


Figure 5: Geographic segmentation (spatial restriction): change in incumbent driver income as a function of the autonomous-vehicle share of trips, when the autonomous vehicle is restricted to the trips with the lowest forward-looking continuation value (FHV 2025, 10-bin $\mathcal{F}$, Section 4.1.2, $\eta = 0.265$). The star marks the peak: +3.2% at an AV share of 4.7% (about 509 AVs). The detailed multi-panel profile (fleet, share, earnings, V_{ante}) and the 2012-baseline counterpart are in Appendix H.

to +1.4% at three-quarters density and to zero by half density, where the labor-optimized deployment is empty. The policy is thus robust to the modeling channels but markedly density-dependent: it needs a market thick enough to leave a clearly low-value tail of trips to hand off.

5.2.6 Implementing the Spatial Restriction

A regulator who accepts that AV entry should not lower incumbent drivers' realized earnings has, in our model, a non-trivial capacity of permitted AV trip-hours. We close the spatial program by describing how that capacity maps to actual regulation.

Historical Precedent: Boro Taxis In August 2013 the NYC Taxi and Limousine Commission created the "Boro Taxis", a green-painted class of licensed street-hail livery cabs barred from picking up in the Manhattan core and at the airports, zones reserved for the existing yellow medallions, with the goal of extending outer-borough service without competing away medallion rents; the class grew to about 6,500 active cabs by 2015 (Mangrum and Molnar, 2020). The same instruments we consider here

(zone designations, permit classes, hail-zone enforcement) could define an AV permit class with a destination-based rule. One difference works in the incumbents' favor: where the 2013 zoning forces green-cab drivers to deadhead home after dropping in dense Manhattan, here the AV fleets absorb the deadhead trips. Human drivers, spared those trips, are simply more productive per active hour.

Congestion concerns The deadheading autonomous vehicles still occupy the road, raising the total number of cars. We include this as a congestion channel, using the density elasticity ($\beta_{\text{cong}} = 0.358$, estimated from the Boro Taxi roll-out itself in Mangrum and Molnar (2020)); while it tightens the permitted fleet modestly, it does not overturn the earnings gain.

What the regulator sets. The instrument here is geographic, defining which trips autonomous vehicles may serve. Fares are regulated and common across yellow, ride-hail, and autonomous vehicles, as under the TLC, and AV wait is fixed, so neither price nor wait is an instrument in this program; the price lever is the subject of Program 1 (Section 5.1). The simplest rule, in the 2025 ride-hail market, is to permit autonomous vehicles only on trips terminating outside the city, which at observed volumes admits about 2,900 vehicles serving roughly 27,000 exterior dropoffs a day; in a higher-friction market the analogous rule is a threshold on the destination's forward value. A restriction of this kind remains useful to AV operators, who can run a fleet and refine their systems on a slice of demand that incumbents are better off to cede.

6 Conclusion

Autonomous vehicles substitute for human driving in the aggregate, so their unrestricted entry would reduce incumbent drivers' trip volume and earnings. We have asked whether regulation can permit substantial AV deployment while protecting incumbent drivers, and compared two programs within a dynamic model of driver labor supply estimated on New York City data.

The canonical remedy (taxing AV revenue and rebating it to drivers) generally cannot generate enough revenue to replace drivers' lost income at current AV costs. Because the autonomous vehicle is close to a pure substitute for the human trip,

the tax recovers only the automation's cost margin while each displaced trip costs a driver the full fare, so the optimal tax-financed fleet is zero. An interior optimum appears only once the AV cost share falls below the composite wait-quality elasticity $\eta|\varepsilon_w|$, the regime in which the technology's productivity gain is large enough that compensation is least needed.

A policy that can condition on which trips are automated succeeds where the tax cannot. Directing AVs to the destinations with the lowest forward-looking value (the trips that strand drivers farthest from their next fare) keeps incumbent earnings at or above baseline across a wide range of deployments and raises their ex-ante welfare, even as AVs serve up to a third of trips. Where dispatch is slow (2012 street hail), human service concentrates in the Midtown–Lower-Manhattan core and AVs take the outer boroughs; where dispatch is fast (2025 ride-hail), drivers keep the whole city and AVs are pushed to trips leaving it. The earnings gains are modest in level (from a fraction of a percent to a few percent) and positive; geography protects workers where taxation reduces their income.

Implementation parallels existing practice. A destination-based permit rule, akin to the 2013 Boro Taxi zoning, allows autonomous vehicles to serve trips terminating in the designated slice and not others. Our analysis holds the driver pool fixed and represents rider demand and congestion through reduced-form elasticities drawn from studies of the same market, rather than a fully structural model of platform and rider behavior; these are its principal simplifications.

References

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.

Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020.

Daron Acemoglu, Andrea Manera, and Pascual Restrepo. Does the U.S. tax code favor automation? *Brookings Papers on Economic Activity*, pages 231–300, 2020. Spring.

David H. Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333, 2003.

Patrick M. Bösch, Felix Becker, Henrik Becker, and Kay W. Axhausen. Cost-based analysis of autonomous mobility services. *Transport Policy*, 64:76–91, 2018.

Nicholas Buchholz. Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies*, 89(2):556–591, 2022.

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

Nicholas Buchholz, Matthew Shum, and Haiqing Xu. Rethinking reference dependence: Wage dynamics and optimal taxi labor supply. *Journal of Political Economy*, 2026. Forthcoming.

Juan Camilo Castillo. Who benefits from surge pricing? *Econometrica*, 93(5):1811–1854, 2025.

M Keith Chen, Peter E Rossi, Judith A Chevalier, and Emily Oehlsen. The value of flexible work: Evidence from uber drivers. *Journal of political economy*, 127(6): 2735–2794, 2019.

Arnaud Costinot and Iván Werning. Robots, trade, and luddism: A sufficient statistic approach to optimal technology regulation. *Review of Economic Studies*, 90(5): 2261–2291, 2023.

Guillaume R Frechette, Alessandro Lizzeri, and Tobias Salz. Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review*, 109(8): 2954–92, 2019.

Joao Guerreiro, Sergio Rebelo, and Pedro Teles. Should robots be taxed? *Review of Economic Studies*, 89(1):279–311, 2022.

Jonathan V Hall, John J Horton, and Daniel T Knoepfle. Ride-sharing markets re-equilibrate. Technical report, National Bureau of Economic Research, 2023.

Anton Korinek and Joseph E. Stiglitz. Artificial intelligence and its implications for income distribution and unemployment. In Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors, *The Economics of Artificial Intelligence: An Agenda*, pages 349–390. University of Chicago Press, 2019.

Daniel Mangrum and Alejandro Molnar. The marginal congestion of a taxi in new york city. Working paper, Federal Reserve Bank of New York, 2020.

Michael Ostrovsky and Michael Schwarz. Carpooling and the economics of self-driving cars. Working Paper 24349, National Bureau of Economic Research, 2018.

Nicola Rosaia. Competing platforms and transport equilibrium: Evidence from new york city. *Econometrica*, 93(6):2235–2271, 2025.

Danial Salman. End of the road? autonomous vehicles and displacement risk. Working paper, University of Washington, Foster School of Business, 2026.

Uwe Thuemmel. Optimal taxation of robots. *Journal of the European Economic Association*, 21(3):1154–1190, 2023.

Waymo. Waymo safety and performance report, 2024.

A Calibration Details

- **Preference estimation.** We estimate $(\theta_1, \theta_2, \sigma)$ per shift type by maximizing the continue/quit likelihood (3) over the weekday panel (24,248 *am*-shift and 11,304 *pm*-shift driver-shifts, owner-operator and lease pooled), holding fixed the nonparametric $\mathcal{F}$; the model continuation probabilities p_θ come from a value-function iteration of (2) at each candidate θ. Decisions before 180 minutes are dropped and the ridge weight is $\lambda = 5$ on σ (income scale, below). The optimum is found by a bounded Nelder–Mead search. Estimates are stored in `data/theta.csv` at the computational income scale; reported values rescale income's coefficient to one (multiplying each θ by ten). Row 1 θ_1 (positive), Row 2 θ_2 (negative), Row 3 σ. Standard errors are obtained by a shift-level bootstrap that resamples whole shifts and re-estimates both $\mathcal{F}$ and θ.
- FHV zone-specific idle: Manhattan 13.5–16.7 min, outer 14.9–16.7 min, Airports 16.7 min, Exterior 41.4 min.
- dh-stretch $\eta = 0.265$: authors' calculation from a 2012 shift-change-window regression of log search time on log active drivers, with day-of-week × month fixed effects.
- AV cost: operating \$0.55/mile (12 mph, 12 hr/day duty cycle, 45% deadhead) plus capital \$100,000/AV amortized over five years give \$72,560/AV/yr and 6.6 loaded trip-hours/AV/day; cost levels follow Bösch et al. (2018).

A.1 Search-time elasticity η: robustness and identification

The default $\eta = 0.265$ is from the 2012 shift-change regression; for FHV it is plausibly an overstatement (modern dispatch is more efficient than street-hail cruising) or an understatement (no spatial slack from rider walk-to-corner substitution). To bound the result, we run two alternative profiles on the 10-bin construction:

- **Lower bound (`data_lower`):** per-zone η set to the FHV PAX-wait regression coefficient $\bar{\beta}$ on 2025 Mon–Thu records (Lower Manh. $\bar{\beta} = 0.11$; Midtown, Upper Manh., Queens $\bar{\beta} = 0.19$; Brooklyn $\bar{\beta} = 0.17$; Airports $\bar{\beta} = 0.13$; pooled 0.16). Bronx and Other-SI assumed at 0.18 (the outer-borough pooled value).

- **Upper bound (`data_upper`):** per-zone η set to $\bar{\beta}/(1-\phi_{\text{supply}})$ with $\phi_{\text{supply}} = 0.6$ (modest within-day supply restaffing). This raises each zone's η to 0.28–0.49.

At the earnings-max deployment across all three profiles:

profile	$B/B_{\max}$	σ_{AV}	Δ earnings
`data_lower` ($\eta = \bar{\beta}$)	0.200	0.154	+3.07%
`yellow_uniform` ($\eta = 0.265$)	0.100	0.047	+2.54%
`data_upper` ($\eta = \bar{\beta}/0.4$)	0.100	0.047	+2.24%

The earnings-max policy is positive in all three profiles. The earnings-positive band shrinks (from $\sigma \leq 0.56$ at the lower bound to $\sigma \leq 0.21$ at the upper bound) but never disappears.

Identification of η from FHV data. The PAX wait-time regression $\log W_{z,h,d,t} = \bar{\beta} \log N_{z,h,d,t} + \text{FE}(z,h,d) + \varepsilon$ identifies the elasticity of consumer wait time to demand variation conditional on zone $\times$ hour $\times$ dow fixed effects. In a constant-elasticity matching model with supply elasticity to demand ϕ_{supply}, the observed $\bar{\beta} = \eta(1-\phi_{\text{supply}})$, so $\bar{\beta}$ is a lower bound on η. For permanent AV deployment, supply restaffing is more limited than for daily demand fluctuations; the relevant η is closer to the upper bound. Direct identification of η would require a supply-side natural experiment (e.g., the August 2018 NYC driver-cap policy as a within-FHV shock) and is a planned extension.

B Out-of-sample validation: the September 2012 fare increase

Our counterfactuals transform $\mathcal{F}$ and re-solve labor supply with θ held fixed; the September 2012 NYC fare increase is the one real shock on which to test that pipeline. We observe the earnings process both before (July–August, our estimation window) and after (September–October) the increase. We take the post-period transition function as directly estimated from the post data and ask whether the stopping model, carrying the *pre*-period preferences $\hat{\theta}$, reproduces the observed pre→post change in behavior. Because the model is fed the realized post environment rather than any attribution of what moved it, the test is robust both to the demand and wait-time

changes the fare increase induces and to September–October seasonality: it asks only whether drivers respond to a changed environment as the model predicts. It therefore validates the labor-supply half of the model (preference stability and the transportability of the stopping rule) and not the separately calibrated $\mathcal{F}$-construction channels.

Table 6 reports the comparison. The roughly 16% fare increase raised earnings per shift and per active hour by about 22% while leaving shift length essentially unchanged, and the model (disciplined only on pre-period quit decisions) reproduces this out of sample (hourly earnings +\$6.9 predicted vs +\$6.6 observed; shift length flat in both). Drivers do not lengthen shifts despite the higher wage, the near-zero intensive-margin response the dynamic model is built to capture. Re-estimating θ on the post period returns similar values (*am*-shift $\theta_1 = 6.7$ vs 7.5; σ and θ_2 within a few percent), consistent with regime-invariant preferences.

Table 6: Fare-increase validation: observed versus model-predicted driver behavior, July–August (pre) and September–October (post) 2012. The model carries the pre-period preferences $\hat{\theta}$ and the directly estimated post-period $\mathcal{F}$; all four moments are untargeted.

	Pre		Post	
Moment	Data	Model	Data	Model
Shift length (hours)	8.1	8.0	8.1	8.2
Earnings per shift (\$)	245	233	300	296
Earnings per active hour (\$)	30.3	29.1	36.9	36.0
Trips per shift	18.2	20.4	18.6	21.8

C 10-Bin $\mathcal{F}$ Construction: Ride-hail Spatial Structure + 2012 Fatigue Profile

This appendix repeats the mathematical formulation of the 10-bin $\mathcal{F}$ introduced in Section 4.1.2, expands the special-case treatment, and notes a planned diagnostic test.

The construction exploits a separation: the *fatigue* component of the 2012 baseline's transition function is a function of cumulative shift time h alone, not of money r or origin location ℓ. Under this assumption, we use the 2012 baseline only for its temporal-fatigue structure and combine it with FHV's own spatial primitives at the

full 10-bin geometry (the seven 2012 NYC regions plus Bronx, Other-Staten-Island, and Exterior, in which FHV does substantial business but the 2012 baseline does not).

State and factorization. Let $s = (r, h, \ell)$ denote a driver's state: cumulative income r, cumulative shift hours h, current location ℓ. A trip from s at location ℓ to next-state location ℓ' produces increments $(\Delta r, \Delta h)$ where Δr is the trip's driver payment and Δh is the trip duration plus post-dropoff idle. We factor the transition function as

$$\mathcal{F}(s, s') \;=\; P(\ell' \mid \ell) \cdot G(\Delta r, \Delta h \mid \ell, \ell', h)\,.$$

The first factor captures spatial movement; the second captures the trip economics including any fatigue effect through h.

FHV primitives (10 bins). For each (ℓ, ℓ') pair with $\ell, \ell' \in \{1, \ldots, 10\}$, FHV's public trip records (2025 Mon–Thu, our estimation window) give:

- $P_{\text{FHV}}(\ell' \mid \ell)$: empirical pickup-to-dropoff probability (location transition).
- $\overline{\Delta r}_{\text{FHV}}(\ell, \ell')$: mean driver pay per trip in the cell.
- $\overline{\Delta h}_{\text{baseline}}(\ell, \ell')$: mean trip duration plus dispatch idle in the cell, measured for "fresh" drivers (i.e., before any fatigue adjustment).

2012 fatigue profile. From the 2012 street-hail market's empirical transition function $\mathcal{F}_{2012}$, estimate the fatigue multiplier

$$\phi(h) \;\equiv\; \frac{\mathbb{E}_{\mathcal{F}_{2012}}\Big[\Delta h \mid h_1 = h\Big]}{\mathbb{E}_{\mathcal{F}_{2012}}\Big[\Delta h \mid h_1 = 0\Big]}, \qquad \phi(0) = 1, \tag{7}$$

where expectations are taken over (origin, destination) pairs present in the 2012 panel. $\phi(h)$ is a monotone-nondecreasing scalar function. The shape of ϕ in 2012 NYC data (a 1.5–1.8× stretch over the first 8 shift-hours) reflects the tendency of tired drivers to search less aggressively post-dropoff.

New $\mathcal{F}$ at 10 bins. For each (s, s') pair with $s = (r_1, h_1, \ell_1)$ and $s' = (r_2, h_2, \ell_2)$, define

$$\mathcal{F}_{\text{new}}(s, s') \;=\; P_{\text{FHV}}(\ell_2 \mid \ell_1) \times \mathcal{B}\Big(r_2 - r_1 \;\approx\; \overline{\Delta r}_{\text{FHV}}(\ell_1, \ell_2)\Big) \times \mathcal{B}\Big(h_2 - h_1 \;\approx\; \overline{\Delta h}_{\text{baseline}}(\ell_1, \ell_2)\,\phi(h_1)\Big), \tag{8}$$

where $\mathcal{B}(\cdot)$ denotes bilinear-stochastic placement on the (i, j) discrete grid of the state space. The interpretation: location transitions come from FHV's empirical marginals; per-trip income is set by FHV's empirical mean; per-trip cumulative-time movement is set by FHV's empirical baseline scaled by the 2012 baseline's fatigue profile $\phi(h_1)$ evaluated at the driver's current shift-time.

Special cases.

- *Exterior outbound* ($\ell_1 \in$ NYC, $\ell_2 =$ Ext): $\overline{\Delta r}_{\text{FHV}}$ and $\overline{\Delta h}_{\text{baseline}}$ come from FHV's empirical Exterior outbound; fatigue $\phi(h_1)$ applies as in (8).
- *Exterior return* ($\ell_1 =$ Ext, $\ell_2 \in$ NYC): $\overline{\Delta r} = 0$ (unpaid), $\overline{\Delta h}_{\text{baseline}} = 41.4$ min (mechanical return-deadhead, from the exterior-row diagnostic). ϕ does not apply because the return is a mechanical infrastructure cost, not a search-driven one.
- *Bronx and Other-SI*: P_{FHV}, $\overline{\Delta r}_{\text{FHV}}$, $\overline{\Delta h}_{\text{baseline}}$ are taken directly from FHV's own data for those bins. The fatigue multiplier ϕ from (7) is applied as in any other NYC bin (assumption: fatigue is location-invariant).

Identification check. The validity of (8) rests on the assumption that the 2012 baseline's $\phi(h)$ transfers to ride-hail. A diagnostic test is available once we have the 10-bin $\mathcal{F}$: re-estimate $\phi(h)$ directly from FHV's trip records (averaging Δh across cells, by binned h_1) and compare to the 2012 baseline's profile. The FHV data lack driver IDs in the public records, so this requires either (i) a synthetic-cohort approximation that imputes h_1 from within-day trip sequences in the same vehicle's data, or (ii) TLC special data access for proper driver-ID linkage. We flag this as a follow-up.

D Demand and Congestion Channels: Implementation Details

Both channels enter the model as multipliers on Δh (the per-trip time increment) at each destination zone in the transition function. They share the same implementation pathway (modifying the per-cell dh_alpha$[o, d]$ matrix and applying via `rescale_wages_cellwise`) but they push in opposite directions. This appendix details the per-zone arithmetic and explains the implementation choice that preserves the BSX continue/quit semantics.

Demand response (FLS)

Economic logic. AV deployment shifts the effective rider-side wait time. Rider demand responds with elasticity $\varepsilon_w = -1.225$ (Frechette et al., 2019). More riders per driver means shorter idle per trip, so Δh shrinks and drivers fit more trips into a shift.

Per-zone arithmetic. At each origin (next-pickup) zone ℓ, with $\sigma_{AV}[\ell]$ origin-conditional and $w_{AV} = 5$ min exogenous AV wait:

1. Post-AV human wait (with dh-stretch already applied):

$$w^{AV}_{\text{human}}[\ell] = w_{\text{human}}[\ell] \cdot (1 - \sigma_{AV}[\ell])^{-\eta[\ell]}.$$

2. Effective rider-side wait (AV serves σ at w_{AV}, humans serve $(1 - \sigma)$ at their stretched wait):

$$\bar{w}[\ell] = \sigma_{AV}[\ell] \cdot w_{AV} + (1 - \sigma_{AV}[\ell]) \cdot w^{AV}_{\text{human}}[\ell].$$

3. Demand scale per FLS:

$$d_{\text{scale}}[\ell] = \left(\bar{w}[\ell] / w_{\text{human}}[\ell]\right)^{-1.225}.$$

In our runs, d_{scale} ranges from 1.006 (very small AV fleet) to 1.80 (largest fleet at $B/B_{\max} = 0.60$).

Feeding into the BSX model. The matching-tightness at zone ℓ after demand

response becomes

$$\text{tightness}[\ell] = (1 - \sigma_{AV}[\ell]) \cdot d_{\text{scale}}[\ell]$$

(instead of just $(1 - \sigma_{AV}[\ell])$ in the no-demand case). The dh-stretch factor at destination ℓ is $\text{stretch}[\ell] = \text{tightness}[\ell]^{-\eta[\ell]}$. When $d_{\text{scale}} > 1$, tightness is larger and the stretch is smaller: idle shrinks because more riders are coming through. The shortened Δh means drivers do more trips per shift.

Why not row-mass scaling. An earlier attempt scaled $\mathcal{F}$'s row mass per origin by $d_{\text{scale}}[\ell]$ to represent "more trips per origin." That approach pushed row sums above 1.0 and broke the BSX continue/quit semantics: in the BSX framework the row residual $1 - \sum_{s'} \mathcal{F}(s' \mid s)$ is the probability of quitting at state s, and with row sums above 1 this residual goes to zero so drivers never quit voluntarily. The dh-shortening approach captures the same economic content (more matches $\Rightarrow$ more trips per shift) while preserving $\mathcal{F}$ as a proper sub-stochastic transition function.

Congestion (Mangrum-Molnar)

Economic logic. AVs added to roads increase vehicle density. Higher density slows travel for all road users (Mangrum and Molnar, 2020). Trip times grow; per-hour earnings fall.

Per-zone arithmetic. At each destination zone d:

1. AV vehicle-hours through zone d:

$$\text{AV-veh-hr}[d] = \sum_o \text{mass}[o, d] \cdot \text{served_frac}[o, d] \cdot \text{trip_time}(o, d).$$

2. Density change as a fraction of baseline:

$$\Delta v/v[d] = \text{AV-veh-hr}[d]/\text{baseline-trip-hours}[d].$$

3. Multiplicative time stretch:

$$\text{cong_factor}[d] = 1 + \beta_{\text{cong}} \cdot \Delta v/v[d], \qquad \beta_{\text{cong}} = 0.358.$$

$\beta_{\text{cong}} = 0.358$ is the Mangrum-Molnar midtown elasticity; the broader treatment-area elasticity is 0.141. We use the midtown value as the conservative (frontier-tightening)

case; the broader value is a less restrictive robustness check. In our runs, cong_factor ranges from 1.008 at small fleets to 1.20 at the largest fleet.

Combined dh-multiplier

The per-cell dh_alpha used in `rescale_wages_cellwise` becomes

$$\text{dh_alpha}[o,d] = \underbrace{\left(1 + \text{search_share} \cdot \Big[\text{stretch}[d] - 1\Big]\right)}_{\text{dh-stretch on idle, demand-aware}} \cdot \underbrace{\text{cong_factor}[d]}_{\text{congestion on trip-time}},$$

with search_share = 0.51. The first multiplier captures matching-tightness effects on the idle component (which interact: AV substitution and demand response both feed through the tightness term); the second captures congestion effects on the full trip-time component. Both human-served and AV-served trips slow by the same factor; both face the same matching-tightness; both feed into the dh-rescale on the deployed $\mathcal{F}$.

Why the channels roughly cancel at the labor-optimized peak

At $B/B_{\max} = 0.10$ ($\sigma_{AV} = 0.05$): the demand-response shrinks the idle stretch from 1.014 (baseline) to 1.007, a 0.7% reduction. This applies to the 51% of dh that is idle, giving a 0.35% reduction in dh. Congestion stretches the trip-time component by 3%, applied to the 49% of dh that is trip-time, giving a 1.5% increase in dh. Net dh growth is +1.15%, a small positive stretch that leaves the peak earnings change in the same 3.0%–3.4% band as the baseline.

At larger fleets the demand-response feedback through $d_{\text{scale}}^{-\eta}$ compounds in the tightness, so demand dominates beyond $B/B_{\max} = 0.3$, which is why the demand-only case stays earnings-positive across the full range of deployments while the congestion-only case turns negative at $B/B_{\max} = 0.30$.

E Robustness: Density and the Demand and Congestion Channels

This appendix details the robustness exercises summarized in Section 5.2.5. The search-time elasticity η is treated separately in Appendix A.1, and the per-zone arithmetic of the two channels in Appendix D.

Density

We re-solve the 10-bin FHV case at two reduced-density scenarios. The construction is the same as Section 4.1.2 but with the dispatch-wait component of per-cell idle scaled by $1/\rho$ for density $\rho < 1$. For NYC destinations this scales idle directly; for Exterior destinations the mechanical 36.4 min return-deadhead is unchanged but the 5-min re-entry wait scales by $1/\rho$. Trip times and the fatigue profile $\phi(h)$ are unchanged.

- $\rho = 0.75$**.** Baseline per-shift earnings fall from \$174 to \$159; hours from 6.24 to 5.72; per-active-hour rate unchanged at \$28. Earnings-max policy at $B/B_{\max} = 0.10$, fleet 573, $\sigma = 0.052$, Δearn $= +1.35\%$ (with demand-response and congestion channels). The earnings-positive band is $B/B_{\max} \in [0, 0.15]$, with maximum fleet at 860 AVs.
- $\rho = 0.5$**.** Baseline per-shift earnings fall from \$174 to \$140; hours from 6.24 to 5.03. Earnings-max policy collapses to $B = 0$: the labor-optimized AV deployment is empty. Every positive deployment delivers negative earnings, -0.24% already at $B/B_{\max} = 0.025$ and -0.62% at $B/B_{\max} = 0.10$.[8]

The earnings gain at the labor-optimized policy scales nonlinearly in density: $+2.57\%$ at $\rho = 1$, $+1.35\%$ at $\rho = 0.75$, and ≤ 0 at $\rho = 0.5$, with the contiguous earnings-positive band shrinking from $[0, 0.40]$ to $[0, 0.15]$ and then closing entirely. The frontier collapses somewhere between $\rho = 0.5$ and $\rho = 0.75$, indicating that the regulatory scope of AV entry is density-sensitive. The mechanism is that sparser markets compress cell-level dispersion: every cell carries longer idle, so the relative gap between best and worst cells shrinks in proportional terms, and the benefit from substituting

[8]The $\rho = 0.5$ run applies the density thinning only; the demand-response and congestion channels are omitted, as they alter earnings materially only at large fleets, where V_{ante} is already negative.



Figure 6: FHV earnings (top) and V_{ante} (bottom) across $B/B_{\max}$ at three market densities, 100% (baseline), 75%, and 50%, from the 10-bin Bellman with $\eta = 0.265$. Gold stars mark the earnings-max policy at each density. The contiguous earnings-positive band is widest at 100% density, much narrower at 75%, and empty at 50%.

low-value trips shrinks accordingly. Meanwhile the dh-stretch on remaining cells scales up because the base idle is larger. The two effects together shrink the earnings-positive band faster than density falls; at $\rho = 0.5$ the band has closed entirely. NYC-based welfare estimates therefore do not extrapolate linearly to smaller metros.

Demand and congestion

The baseline FHV exercise holds rider demand and road congestion fixed. Two complementary channels are missing from that specification: **demand response** (Frechette et al., 2019) (when AV shortens effective rider wait time, total demand

grows, partially offsetting the displacement of human-served trips) and **congestion** (Mangrum and Molnar, 2020) (AVs on the road slow all travel, tightening the labor-optimized frontier). We re-solve the 10-bin Bellman under each channel separately and then together; the per-zone arithmetic of both channels is in Appendix D.

At the labor-optimized deployment $B/B_{\max} = 0.10$ ($\sigma_{AV} = 0.047$, fleet 509):

- **Baseline:** Δearn $= +2.54\%$, $\Delta V_{\text{ante}} = +3.87\%$. Earnings-positive band $[0, 0.35]$.
- **Demand only:** Δearn $= +2.72\%$ at $B/B_{\max} = 0.10$. Earnings stay positive across the entire deployment range $[0, 0.60]$ (the global peak of $+3.37\%$ falls at $B/B_{\max} = 0.525$, but V_{ante} is negative there). The frontier shifts outward.
- **Congestion only:** Δearn $= +2.39\%$ at $B/B_{\max} = 0.10$. Earnings turn negative at $B/B_{\max} = 0.30$, shrinking the earnings-positive band to $[0, 0.25]$. The frontier shifts inward.
- **Both channels:** Δearn $= +2.57\%$ at $B/B_{\max} = 0.10$, essentially unchanged from baseline at the peak. The two channels roughly cancel at the labor-optimized policy point. At larger fleets demand response dominates: earnings dip near the $B/B_{\max} = 0.45$ cell-threshold and re-emerge positive (peak $+3.03\%$ at $B/B_{\max} = 0.525$), though V_{ante} is negative across that region.

Demand response expands the set of welfare-positive deployments and congestion shrinks it, leaving the labor-optimized policy point unchanged.

F 2012-Baseline V_{ante} vs. Realized Earnings

- $V_{\text{ante}} \approx \mathbb{E}[U(r_{\text{final}}, h_{\text{final}})] + \mathbb{E}[N]\sigma\gamma$.
- With placeholder $\sigma \approx 6.5$–7 in theta-units (\$65–\$70 dollar Gumbel scale), $\sigma\gamma \approx$ \$38/decision; $\mathbb{E}[N]\sigma\gamma \approx$ \$650/shift.
- Channels in the $+10.6\%$ V_{ante} gain: realized income ($+5\%$-ish), realized hours (slight negative under hump-shaped utility), $\sigma\gamma N$ (positive), and a residual we plan to decompose with a simulator diagnostic dump.
- Practical point for the paper: V_{ante} and earnings peak at the same $B/B_{\max}$ in the 2012 baseline, so the policy choice is robust to which metric one commits to.

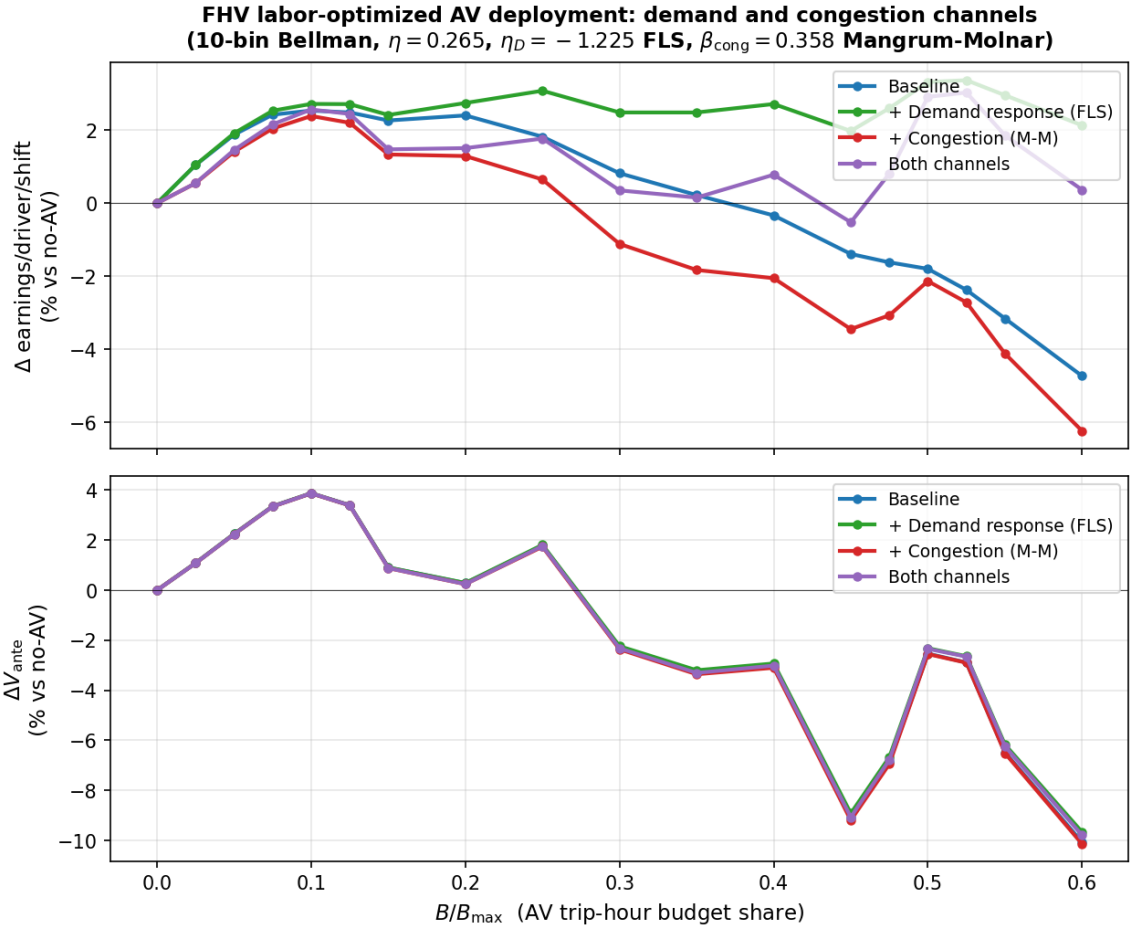


Figure 7: FHV earnings (top) and V_{ante} (bottom) under four specifications: baseline (no channels), demand response only ($\varepsilon_w = -1.225$), congestion only ($\beta_{\text{cong}} = 0.358$), and both. All four use the 10-bin Bellman with $\eta = 0.265$. Demand response shifts the earnings-positive band outward; congestion shrinks it.

G Modality Comparison Figure

- 2012 yellow: 91% Manhattan-internal pickups.
- 2025 yellow: 84% Manhattan-internal pickups.
- 2025 FHV-HV: 28% Manhattan-internal; Brooklyn alone 27%; Bronx, Staten Island, Newark populated.
- Within-NYC trip-flow imbalance among 232 non-trivial zones: largest is ±1 percentage point, volume-weighted standard deviation within Brooklyn / Queens / Bronx below 0.03 percentage points.

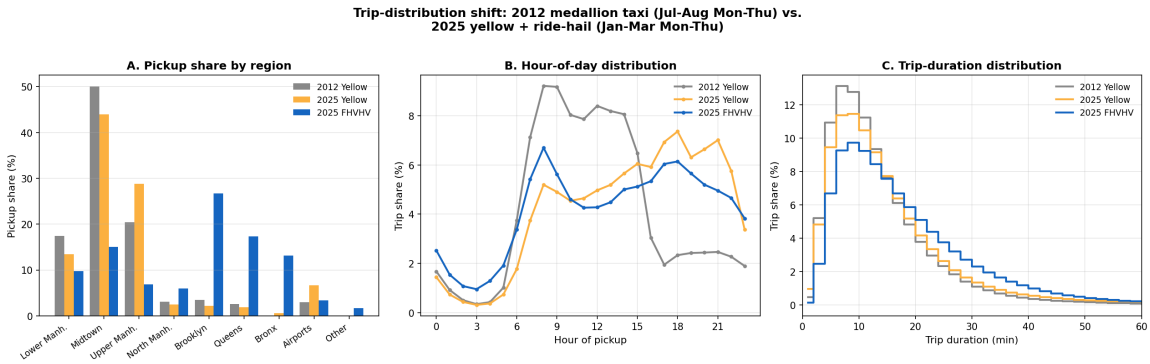


Figure 8: Trip-distribution shift, 2012 street-hail vs. 2025 yellow and FHV-HV (Mon–Thursday).

- Implication: modern FHV dispatch compresses within-NYC dispersion, so the destination-based AV slice is the network exterior. In 2012 the 2012 baseline, by contrast, dispatch friction puts the boundary inside NYC itself, between the Midtown / Lower-Manh. core and everything else.

H Detailed AV Deployment Profiles

Figure 5 in the main text reports incumbent driver income as a function of the autonomous-vehicle trip share for the FHV 2025 segmentation. This appendix gives the full multi-panel profiles underlying that figure, for both the 2025 ride-hail and the 2012 markets.

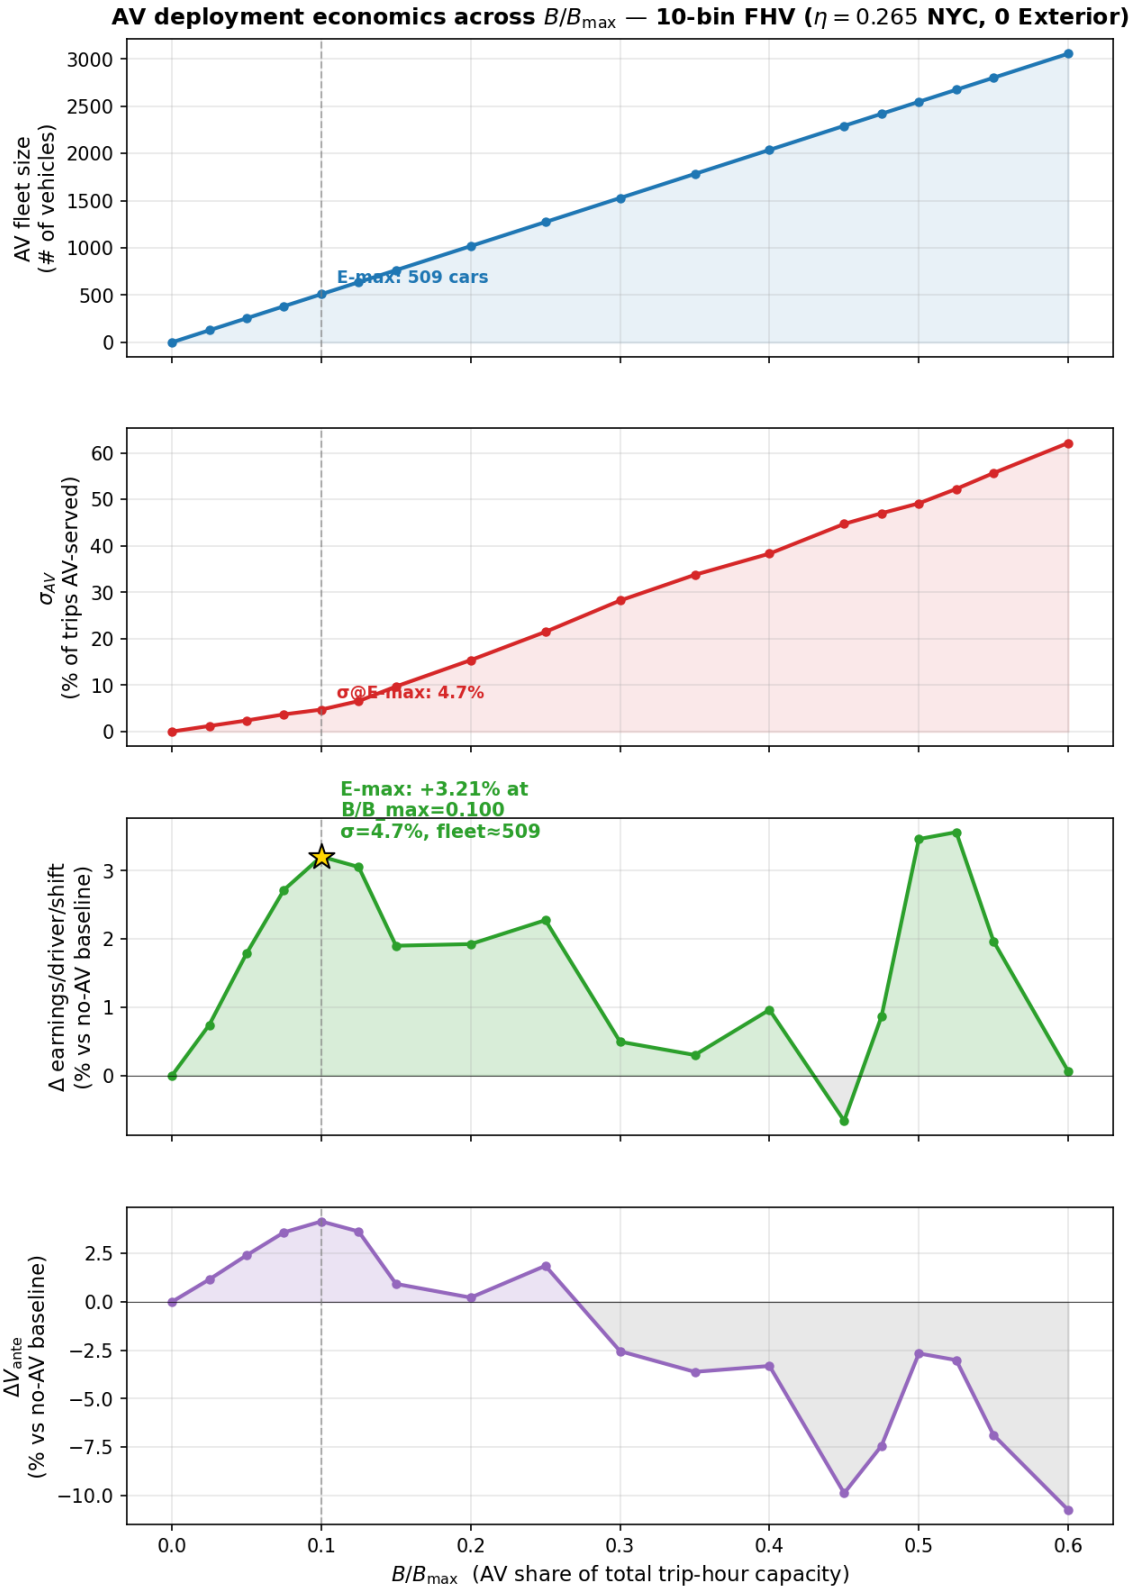


Figure 9: FHV 2025 deployment profile across the AV capacity $B/B_{\max}$, from the BSX simulator on the 10-bin FHV $\mathcal{F}$ (Section 4.1.2) under the default $\eta = 0.265$ profile. Panels: AV fleet; AV trip share σ_{AV}; change in per-shift earnings; change in V_{ante}. Per-shift earnings rise across $B/B_{\max} \in [0.025, 0.375]$, peaking at +3.2% at $B/B_{\max} = 0.10$ (fleet 509 AVs, $\sigma_{AV} = 4.7\%$); the earnings-positive band ends near $B/B_{\max} = 0.45$.

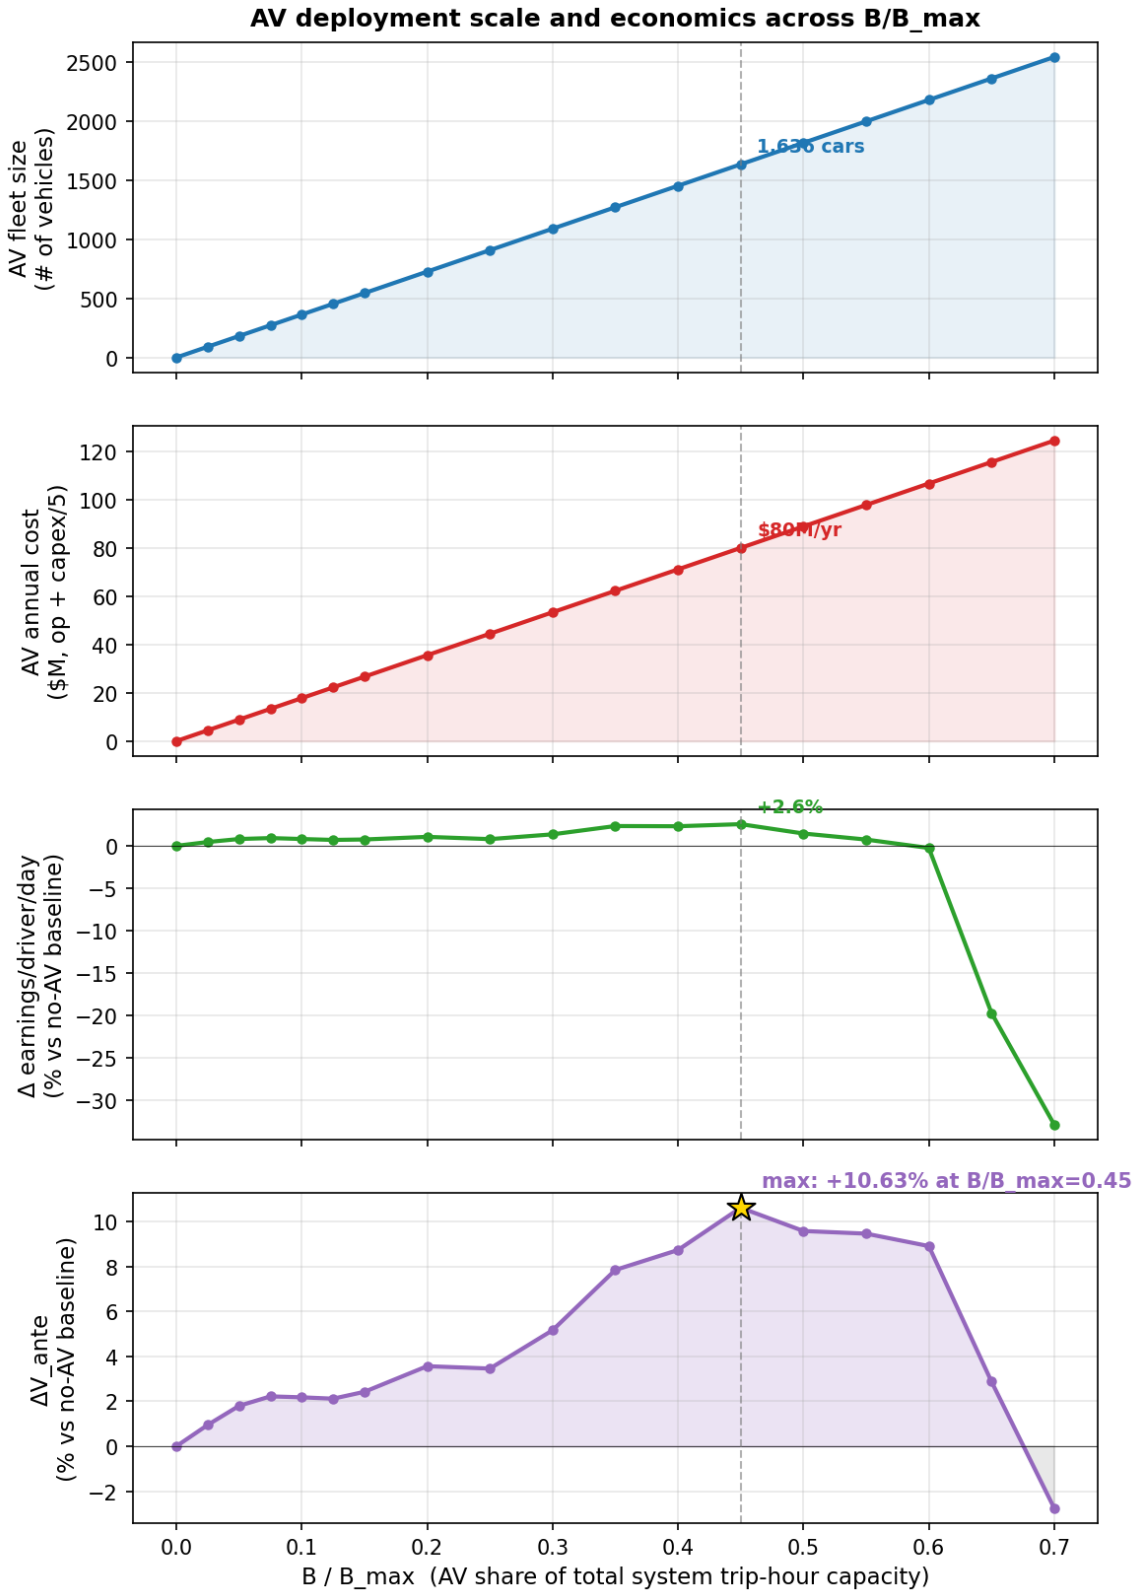


Figure 10: 2012-baseline deployment profile, four panels (fleet, cost, change in per-shift earnings, change in V_{ante}) from the BSX simulator on the empirical 2012 $\mathcal{F}$. Earnings and V_{ante} peak together at $B/B_{\max} = 0.45$.