Will New Driving Technologies Change the Value of Public Transportation Investments?*

Imke Reimers
Cornell University and ZEW

Benjamin Reed Shiller
Brandeis University

June 19, 2025

Abstract

We analyze how self-driving vehicles (SDVs) influence commuter behavior and returns to long-lived public transit investments. Using a commuting mode model estimated on detailed home and work location data from Greater Boston, we simulate the widespread entry of SDVs, which offer passive travel similar to transit but use existing road networks. We find that SDVs increase vehicle miles by 40% while decreasing public transit use by about 10%. Transit improvements continue to moderately boost revenues and lower miles driven, but their effects on mileage are small compared to SDVs. These findings highlight planning challenges posed by the emergence of SDVs.

*We thank Nafisa Lohawala, Adriaan Soetevent, and Joel Waldfogel, as well as conference and seminar participants at Brandeis University and the 2025 International Industrial Organization Conference, for valuable comments and feedback.

1 Introduction

Long-lived transportation infrastructure requires substantial upfront investment, with uncertain returns that hinge on evolving technologies. The emergence of self-driving vehicles (SDVs) may be particularly disruptive: As autonomous driving technology reduces the disutility of travel time, it effectively lowers the opportunity cost of car commuting, potentially altering established patterns of both driving and transit use, as well as residential location choice.[1] In this paper, we examine the interplay of these demand-side innovations and improved public transportation options in commuter choices. Specifically, we ask (1) how autonomous driving technology will affect residential location and commuting mode, (2) how these changes affect the returns to investments in public transit, and (3) whether the public transit improvements can offset the effects of SDVs on driving and transit revenue.

To evaluate the effects that SDVs will have on commuter choices and the returns to investment in public transportation, we first need to quantify consumers' inherent preferences for home locations and commuting options. Using comprehensive data on home and work locations for Boston area workers, along with detailed information on commuting times by transit mode choice, we jointly estimate demand for housing location and transportation choices by income level. Specifically, we model transportation mode choice as a function of its inherent attributes: active driving time, passive travel time (such as on a train), number of stops, and prices. We infer preferences for these characteristics through the individuals' residential location choices given their work locations. The degree to which each transportation mode feature influences home location choice indicates the strength of preference for that feature; and the variation of work locations across our sample allows

[1]For example, Mercedes recently began selling Level 3 autonomous cars with hands-off technology, and it has been testing fully autonomous driving technology (Level 4) in China since August 2024. See https://www.forbes.com/sites/kyleedward/2023/09/28/mercedes-benz-first-to-gain-us-approval-for-level-3-automated-driving-system/ and http://wardsauto.com/autonomous-adas/mercedes-benz-gains-level-4-autonomous-driving-test-license.

us to disentangle preferences for transportation characteristics from intrinsic preferences for specific residential areas. We validate our model by comparing estimated commuter rail ridership with aggregate and location-specific train ridership information from separate data sources.

Our estimates suggest that consumer disutility for each minute of active driving is over 50 percent higher than their disutility of passive commuting time. We interpret these estimates as evidence that consumers value having their attention free for other activities during commuting time. We use these commuting preference estimates to analyze our three main research questions.

First, we examine how SDVs affect consumer commuting choices and quantify the resulting impacts on miles driven and public transit revenue. To model the impact of SDVs, we assume that the disutility associated with time spent driving will be reduced to that of passive travel on public transit, while maintaining the original driving time subject to congestion changes. Holding fixed the public transit network, we find that the widespread adoption of SDVs raises miles driven by between 33 and 41 percent, depending on the housing supply elasticity. This increase arises from three separate factors: (1) the direct positive effect of substitution toward driving, (2) a positive indirect effect of individuals moving further away due to lower driving disutility, and (3) a negative indirect effect of higher congestion lengthening driving times conditional on distance. In addition to raising driving miles, the replacement of conventional cars with SDVs also reduces public transit revenue, by 13 percent; and the share of individuals using public transit falls from 14.3 percent to 13.3 percent. These shifts in commuting behavior are accompanied by meaningful gains in consumer surplus. Partial transition from conventional vehicles to SDVs produces effects that are proportional to the rate of adoption, and the rise in vehicle miles traveled remains substantial even under substantially higher parking prices.

While the arrival of SDVs decreases public transit use overall, the answer to our second question, whether SDVs change the marginal returns to investment in public transit improvements, requires further analysis. We examine these impacts by simulating transit improvements both in scenarios with and without SDVs, and we find that such improvements reduce miles driven and increase day-to-day transit revenue by similar magnitudes regardless of SDV adoption.

Finally, we ask whether improvements in public transit can counteract the effects of SDVs on driving behavior and transit revenue. We find that even substantial enhancements to transit only have modest effects on car mileage. For example, a 40 percent reduction in on-board travel time reduces revenue by just 3 percent. With this transit improvement, our estimates suggest that any SDV adoption rate above 8 percent would still raise driving mileage relative to the status quo. By contrast, public transit improvements can generate sufficient transit revenue to offset the losses caused by SDVs. Even a 20 percent speed improvement raises transit revenue well above status quo levels without transit improvements or SDVs. Our results indicate that improvements in public transit can still raise revenues and yield material benefits for riders, even if SDVs substantially raise miles driven.

The remainder of the paper proceeds as follows. In Section 2, we provide background on private and public transportation options, including the development and take-up of autonomous driving technologies. We also discuss the relevant literature. Section 3 describes the two main data sources used in this paper, on home and work locations and on commuting and transit times between those locations. Section 4 then introduces the model and empirical estimation strategy. The section also highlights the main identification strategy and validates our model by comparing predictions with auxiliary datasets. In Section 5, we use the estimates from the model to evaluate counterfactual scenarios with SDVs as well as with improvements in public transportation. Finally, Section 6 concludes with implications for current investment decisions and long-term policies.

2 Background

This paper's examination of two separate changes in commuting options—SDVs and public transit enhancements—is motivated by recent developments in transportation technology and ongoing investments in public transit. In this section, we lay out the current state of public transit and driving technologies, and we describe how our paper connects to recent related literature.

2.1 Commuting Choices

Individuals living in one location and working in another can choose among various options for commuting. In the United States, two of the most commonly used options include driving and—where possible—using public transportation. Driving continues to be the top choice with about 72 percent of commuters reporting using their car to travel to work in 2024.[2] Public transportation, at about 14 percent, is the second choice, just ahead of walking (12 percent) and riding a bicycle (9 percent). Transportation choices may be influenced by the convenience of each mode, which in turn depends on investments in infrastructure and technological innovations.

2.2 Public Transit Investments

Public transportation takes on a sizable—and growing—share of total government spending across the US. The congressional budget office reports annual investments for public transit

[2] See https://www.statista.com/forecasts/997176.

projects from 1956 to 2014.[3] Over this time period, the share of total transportation infrastructure spending that is allocated to mass transit increased from about seven percent in the 1950s and 1960s to over a quarter in the 2010s.[4] In absolute terms, total nominal US spending on mass transit has increased from $580 million in 1956 to over $65 billion in 2014, a 13-fold increase in real terms. Even more pronounced, capital investment for mass transit has increased from $90 million to $22.7 billion over the same time span, amounting to a 29-fold increase in real dollars.

The city of interest in our study, Boston, also bears high costs of running and updating its public transit system. The Massachusetts Bay Transportation Authority (MBTA) runs four types of public transportation modes: commuter rails, which connect the outskirts of Boston to the city; boats, which connect locations near the coast; trams to get around locations closer to the city center; and buses offering routes that are not serviced by other means. The systems require estimated expenditures of about $25 billion to address aging tracks, stations, and other infrastructure issues;[5] and any enhancements to the infrastructure would increase costs further. For example, speeding up commuter rail lines by electrifying the commuter rail, and reducing wait times between train arrivals, would raise expenditures by another $30 billion or more.[6] Together, these costs would amount to about $8,000 per Massachusetts resident. These high costs raise questions about whether these investments will retain their intended benefits in the presence of SDVs.

[3] See https://www.cbo.gov/sites/default/files/114th-congress-2015-2016/reports/49910-infrastructuresupplementaldata.xlsx.

[4] This rise is remarkable when considering that real spending for the Interstate Highway System rose more than threefold just from the 1960s to the 1980s (Brooks and Liscow, 2023).

[5] See https://www.boston.com/?post_type=post&p=28555994.

[6] See https://cdn.mbta.com/sites/default/files/2019-10/rail-vision-alternative6-oct2019-accessible.pdf. More specific investments are listed in the MBTA's Capital Investment Plan at https://www.mbta.com/financials/capital-investment-plan.

2.3 Self-Driving Vehicles

The history of cars has been defined by periods of technological advancements, including propulsion systems, safety features (such as air bags, electronic stability control, and automatic emergency braking), as well as measures to facilitate driving. The latter group includes the introduction of automatic transmissions and cruise control as first steps toward driverless cars.[7] Many of these technologies followed a familiar diffusion pattern: initial introduction in select vehicles followed by widespread adoption once costs declined sufficiently. Following this pattern, SDVs have seen immense improvements in recent years; and looking forward, they may substantially change how individuals travel between locations.

The Society of Automotive Engineers (SAE) has established a widely recognized framework for classifying the levels of self-driving automation, ranging from Level 0 (no automation) to Level 5 (full, independent automation under all circumstances); and recently introduced cars have moved up to Level 4 on this scale. Tesla's cars enabled with Autopilot or the more advanced "Full Self-Driving" provide a prominent example of Level 2 cars. They allow the system to manage both steering and acceleration/deceleration under certain conditions but require constant driver supervision. One step beyond this, Level 3 offers conditional automation, where the system handles all driving tasks, permitting the driver to disengage (e.g., read a book) but demanding readiness to retake control if notified. In 2023, Mercedes-Benz became the first manufacturer of commercially-available vehicles to be approved by the state of California to offer Level 3 self-driving up to 40 miles per hour on select roads and in favorable conditions (Neil, 2023). Another step further, Waymo, a subsidiary of Alphabet, offers Level 4 autonomous rides without a driver in geographically bounded areas in Austin TX, Los Angeles CA, Phoenix AZ, and San Francisco CA, and it is planning to expand to

[7] See https://opentextbooks.clemson.edu/sts1010fidlerfall2021/chapter/the-evolution-of-automobiles/ and https://www.daleadams.com/blogs/news/20-greatest-innovations-inventions-of-automobile-engineering-from-interestingengineering-com.

Atlanta GA and Miami FL.[8] Finally, no vehicles are currently classified as Level 5.

Advancement in autonomous driving hinges on simultaneous progress across multiple dimensions, including robust sensor technologies (such as lidar, radar, and cameras), sophisticated machine learning models for interpreting sensor data, and high-performance in-car hardware capable of real-time processing. Although current self-driving vehicles demonstrate that the technology is viable, widespread adoption is likely contingent on affordable hardware. One current bottleneck, lidar, remains exclusive to certain luxury vehicles like the Mercedes-Benz EQS, S-Class, and the Volvo EX90. Following the "Varian Rule" on technology adoption patterns, this advanced feature can be expected to reach mid-range vehicles within 5 years and economy models within a decade.[9]

With increasingly affordable autonomous driving technology, SDVs may soon be widely available to the public. Owning such a vehicle will resemble having a private taxi service or personal driver available on a moment's notice (and at a substantially lower per-ride cost). However, the rate of adoption will depend on two key factors: how much consumers value using commute time for other activities, and how strongly they prefer maintaining direct control of their vehicle (e.g., due to safety concerns). Although some are skeptical about SDVs, citing safety issues, the share of consumers who are willing to embrace these technologies has been growing. For example, in a Pew Research Center study from 2017, 39 percent of survey respondents indicated they were "very" or "somewhat enthusiastic" about driverless vehicles, although 53 percent indicated being "very" or "somewhat worried" (Smith and Anderson, 2017). The same study reports that "roughly two-thirds of Americans expect most cars to be driverless in [the] next half century." A 2024 study by AAA notes that skepticism toward SDVs persists, suggesting that investments in safety will be crucial to their success.[10]

[8] See https://waymo.com/, accessed June 13, 2025.

[9] See https://www.theguardian.com/commentisfree/2015/apr/26/facebook-isnt-charity-poor-pay-by-surrendi

[10] See https://newsroom.aaa.com/2024/03/aaa-fear-of-self-driving-cars-persists-as-industry-faces-an-unc

2.4 Related Literature

Our broad question—about the effects of new technologies and public transit investment on home location and transportation choices—relates to three literatures. First, we examine the impacts of technological changes on transportation markets. Recent research has shown that ride hailing services such as Uber can increase public transit ridership by addressing first- and last-mile connectivity issues (Hall et al., 2018). Several papers document non-negligible improvements in efficiency and lower search frictions in these markets (Cramer and Krueger, 2016; Chen et al., 2019; Frechette et al., 2019; Guda and Subramanian, 2019), which can be especially valuable during work hours (Buchholz et al., 2025).[11] However, these benefits are accompanied by increased congestion and environmental costs (Tarduno, 2021), although work-from-home arrangements provide a potential counterbalance to these increased congestion and environmental externalities (Davis et al., 2024; Delventhal and Parkhomenko, 2024; Ioannides et al., 2014; Monte et al., 2023).

While recent technology adoptions such as ride-hailing services are well studied, the scarcity of empirical studies on SDVs reflects their limited deployment to date. Nevertheless, theoretical work and surveys highlight potential impacts. SDVs offer several technological advantages over human drivers: enhanced sensors, faster reaction times, and vehicle-to-vehicle communication capabilities. These features improve safety (Agrawal et al., 2019) and potentially enable higher speeds by reducing vehicle spacing, which can in turn increase road capacity (Hyldmar et al., 2019). However, SDVs may also increase congestion by encouraging longer commutes (Larson and Zhao, 2020).

Second, our analysis connects to research on infrastructure investment and urban transportation policy. Recent work has developed sophisticated frameworks for evaluating trans-

[11] Search frictions have also been studied in more traditional settings. Recent examples include Buchholz (2022) and Brancaccio et al. (2023).

port network improvements in spatial equilibria (Fajgelbaum and Schaal, 2020; Allen and Arkolakis, 2022), with empirical applications showing substantial heterogeneity in returns across different infrastructure segments. Studies of major transit investments find modest but meaningful benefits: Severen (2023) estimates annual benefits of 12 to 25 percent of costs for Metro Rail, while Tsivanidis (2023) shows that standard approaches based on travel time savings capture only about half of the total welfare gains from Bus Rapid Transit. The introduction of SDVs may fundamentally alter these cost-benefit calculations. Our paper examines this possibility.

Finally, our estimation of home location choices builds on literature examining how transportation innovations affect urban form and travel patterns.[12] Existing evidence shows that transportation infrastructure can dramatically reshape cities (Baum-Snow, 2007; Jerch et al., 2024), although road infrastructure or transit improvements need not reduce commuting times due to induced demand and spatial resorting (Van Ommeren and Rietveld, 2005; Duranton and Turner, 2011; Hsu and Zhang, 2014). Several recent papers also account for consumer reoptimization. They examine the equilibrium welfare impacts of various interventions, including transit improvements and network reorganizations (e.g., Barwick et al., 2024; Kreindler et al., 2023), as well as road restrictions and transit pricing (e.g., Almagro et al., 2024; Durrmeyer and Martinez, 2024; Kreindler, 2024). While several analyses focus on market-specific equilibrium effects, some also emphasize the importance of considering environmental (e.g., Glaeser and Kahn, 2010) or speed (e.g., Couture et al., 2018) externalities.[13] We contribute to this literature by exploring how technological advances will reshape location choices and the effectiveness of infrastructure investments.

[12] See Weisbrod and Reno (2009) for an overview.
[13] For a broader overview, see Kahn and Walsh, 2015.

3 Data

The main analysis combines data from three sources. First, we obtain the counts of people living and working in distinct census tracts in the greater Boston area within Massachusetts from the US Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) survey. Second, we combine these data with information on travel times between the two locations for three commuting options: driving, public transportation, and a hybrid option of driving to public transit access points. We supplement these data with commuting and parking prices.[14]

The LEHD Origin-Destination Employment Statistics (LODES) dataset provides information on home and work locations by census tract pair.[15] For each origin/destination pair, the dataset reports the count of workers who live in the "origin" tract and work in the "destination" tract, both in total and by income category.[16] To make the data collection and empirical estimation tractable, we restrict our attention to the 100 work census tracts with the most workers within 20 miles of Boston's city center (City Hall) as well as all 1301 home census tracts within a radius of 75 miles around Boston's city center.[17] Our sampling retains about 35 percent of all workers in the LODES Massachusetts dataset: Restricting work locations to fall within 20 miles of Boston retains 56.6 percent of workers in the dataset, and restricting to the 100 most worked locations retains 61.9 percent of the remaining workers.

To systematically analyze commuting patterns and preferences for different transportation characteristics, we collected travel times between the geographical centers of the home and work census tracts for the three transportation modes.[18] We obtain driving times and

[14] See https://www.boston.com/cars/car-guides/2019/02/24/guide-to-boston-parking-garages/.

[15] See https://lehd.ces.census.gov/data/lodes/. We use version 7, which uses 2010 census block definitions and contains data from 2018.

[16] The dataset covers approximately 95 percent of the U.S. workforce (Graham et al., 2022).

[17] In the model described below, home locations beyond 75 miles from the city center constitute the outside good.

[18] For the hybrid option, we computed times using the closest stop on each commuter rail line and selected

distances from the OpenRouteService API, which accounts for typical congestion levels;[19] and we calculate public transit times using the MBTA Trip Planner.[20] For public transit and hybrid routes, we specify arrival times of 9:00 AM on Wednesdays to reflect typical commuting conditions. We record travel time and the number of transfers between segments.

We supplement the commuting time information with data on one-way transportation prices. Specifically, we use the one-way transit cost from the MBTA Trip Planner for public transit segments. For driving segments, we calculate the commuting cost as $0.18 per mile—which is the IRS-suggested variable cost for 2018—plus half of the daily destination parking cost inferred from the closest public parking lot.[21] Public transit prices seem to be set exogenously without consideration of cross-location variation in demand. Figure 1 plots transit prices for different commuting zones, in 1989 vs. 2024. The prices lie almost exactly on a straight line—with the exception of the lowest-price zone (1A)—suggesting minimal relative price changes over a 35-year period that appear to reflect rounding.

Table 1 shows summary statistics for commuting times, weighted by the number of people living and working in each census tract pair. The top panel shows travel times if using only public transportation.[22] The average public transit commute takes 74.4 minutes (median = 66) and includes 2.2 (2) segments. This is slightly shorter than commutes that combine driving and transit, where the average (median) combined commuting times are 85.5 (77.7), although much of this difference is explained by the fact that public transit-only routes are less likely to be available for the longest commutes. Driving, on the other hand, tends to be much faster: On average, workers can reach their work location by car in 28.9 minutes

the fastest resulting route.

[19] Validation tests on a sample of routes revealed that OpenRouteService's average travel times closely matched Google Maps' rush hour estimates.

[20] See https://www.mbta.com/trip-planner.

[21] The IRS designates the medical care transportation cost per mile as the standard measure of the variable cost of driving. See https://crsreports.congress.gov/product/pdf/IN/IN12320.

[22] The number of observations varies across transportation modes because not all census tract pairs can be connected using public transit.

(median = 23.6). Only 0.2 percent of census tract pairs have faster public transit connections than driving connections, partly reflecting additional time needed to travel to and from stations and to transfer between lines when using public transit (Glaeser et al., 2008).

Despite the fact that transit is rarely the quicker option, a non-negligible share of individuals still chooses public transit, as documented in Section 2.1. This suggests that when travel times are similar between driving and transit, people often prefer using transit. We leverage the variation in the time difference between driving and transit journeys across tract pairs in our main analysis.

Average commuting times and segment counts across all census pairs—not weighting by chosen locations—are substantially longer than those reported in Table 1, at about 101 minutes (vs. 74) and 2.7 segments (vs. 2.2) for public transit connections, and 50 minutes (vs. 29) for driving connections. This difference between raw and weighted commuting times corroborates our initial conjecture that people value shorter and more convenient commutes.

Figure 2 provides additional suggestive evidence to support this conclusion. The figure shows a heatmap of residential densities (of employed residents) along with main components of the transportation network. It indicates some bunching in census tracts near commuter rail lines and major highways. While consumers seem to value shorter commutes, understanding the extent of these preferences, and how these preferences depend on the mode of transportation, requires a model of location and transportation choice. We turn to this in in the next section.

4 Model and Estimation

Our goal is to estimate consumer preferences for home locations as well as for modes of transportation, for both high-income and low-income consumers. We obtain our estimates through a residential choice model that treats work locations as given.[23] The basic premise is that, ceteris paribus, individuals are more likely to reside in areas with appealing transportation options to their workplace. This allows us to infer transportation preferences by examining where they choose to live without needing information on their daily travel decisions. We describe the process and results in this section. In the next section, we will use our estimated preferences to examine both home location and transportation choices when transportation mode characteristics change.

4.1 Model

We begin by specifying the consumer's utility from transportation options, and we then incorporate this utility into their home location preferences. The conditional indirect utility that consumer i in income group g (from set G) receives from transportation mode c (from set $C = \{car, transit, both\}$) when traveling along an origin (h)/destination (w) pair is:

$$u_{icghw} = X_{chw}\beta_g + \epsilon_{ichw}, \tag{1}$$

where X_{chw} denotes the characteristics of transportation option c, β_g describes the group-specific preferences for these characteristics. The outside option, which includes alternatives like bicycling, walking, and working from home, is assigned a mean utility of zero.

Assuming ϵ follows the type 1 extreme value distribution, the expected maximum utility at a given home location stemming from the various transportation options is the inclusive value λ:

$$\lambda(X_{hw}, \beta_g) = E\left[\max_c(u_{icghw})\right] = \ln\left(1 + \sum_{c' \in C|hw} \exp(X_{chw}\beta_g)\right), \tag{2}$$

[23] Barwick et al. (2024) demonstrate through an event study analysis that changes in workplace location typically preceded home purchases, not vice versa, lending credence to the treatment of work locations as exogenous.

where X_{hw} denotes the characteristics of all available transportation options between h and w.

The conditional indirect utility of residential location choice h is then comprised of a home-location specific utility derived from local amenities (δ_{gh}), the utility derived from the commute ($\lambda(X_{hw}, \beta_g)$), and a type 1 extreme value error term (η_{ihw}):

$$\mathbb{u}_{ighw} = \delta_{gh} + \lambda(X_{hw}, \beta_g) + \eta_{ihw}. \tag{3}$$

Then, the probability that a type g consumer who works in location w chooses to live in location h equals:

$$\mathbb{s}_{ghw}(\delta_g, X_{hw}, \beta_g) = \frac{\exp\left(\delta_{gh} + \lambda(X_{hw}, \beta_g)\right)}{1 + \sum_{h' \in H} \exp\left(\delta_{gh'} + \lambda(X_{h'w}, \beta_g)\right)}, \tag{4}$$

where δ_g denotes the set of preferences for all home locations.

Given particular values of β_g, we approximate group g's non-commuting component of utility derived from the home location (δ_{gh}) by iterating on the Berry et al. (1995) contraction:

$$\delta_{gh}^{(n+1)} = \delta_{gh}^{(n)} + \ln(S_{gh}) - \ln\left(\sum_{w \in W} \mathbb{s}_{ghw}(\delta_g^{(n)}, X_{hw}, \beta_g) f_g(w)\right), \tag{5}$$

where S_{gh} denotes the observed overall share of workers of type g that choose to live in location h, and $f_g(w)$ denotes the share of workers of type g employed in location w.

Finally, given β_g and the corresponding δ_{gh}, the log likelihood of workers' home location choices equals:

$$\ln(L(\beta_g)) = \sum_{h \in H, w \in W} \#_{ghw} \ln(\mathbb{s}_{ghw}(\delta_g, X_{hw}, \beta_g)), \tag{6}$$

where $\#_{ghw}$ denotes the observed number of consumers in group g who live at location h

and work at location w.

4.2 Identification and Price Exogeneity

Our goal is to use the model to estimate preferences for home locations and transportation choice characteristics, including transportation mode-specific travel time, costs, and the number of transfers. Here, we explain how preferences for home and transportation choices are identified. We also examine potential confounders of price effects.

4.2.1 Identification

We derive the spatial mean utilities δ_{gh} through the application of a well-established contraction mapping algorithm (Berry et al., 1995). The inferred appeal of each home location encapsulates the aggregate utility derived from location characteristics such as school quality, crime incidence, housing costs, and other local amenities.

Transportation preferences are identified through spatial variation in employment locations. Whereas home location amenities are common across all consumers in group g and location h, transportation attributes depend on both an individual's home and workplace locations. The strength of preferences for specific transportation attributes is identified through the degree to which consumers systematically select residential locations that excel in particular transportation characteristics given their workplace location.

To build intuition, consider the case of an employee at Brandeis University (in Waltham), which is directly served by the Fitchburg commuter rail line. For a worker living in Concord—also on the Fitchburg Line—the public transit commute takes approximately 20 minutes.[24] Driving during rush hour takes about the same time, ∼20 minutes. By contrast, a worker

[24] Travel times from Google Maps for a 9am arrival.

living near Boston Landing would face a similar driving time but a ∼60-minute public transit commute requiring two transfers.

An employee at Wellesley College faces the reverse. Car travel times from Concord and Boston Landing are again comparable, but public transit is far faster from Boston Landing (∼30 minutes) than from Concord (∼150 minutes, two transfers).

Since Brandeis and Wellesley employees face the same housing costs and amenities at each location, variation in residential choices by workplace isolates preferences for specific transportation features. If Brandeis employees disproportionately choose Concord while Wellesley employees prefer Boston Landing, the model infers that consumers value convenient public transit options.

Variation in train speeds provides additional identifying variation. For example, in Boston's subway system, the Orange Line, which runs primarily underground between the northern and southern parts of the city, maintains relatively high speeds averaging 15 miles per hour.[25] By contrast, the Green Line, despite sharing two stops with the Orange Line, moves slower, at an average of only 10 miles per hour; and some segments move as slowly as 3 miles per hour.[26] Higher propensities to live near the orange line, at similar distances from workplaces, indicates that commuters value faster service.

4.2.2 Price Exogeneity

The validity of our empirical analysis also rests on addressing potential endogeneity concerns in both the choice of residential location and transportation costs. In urban choice models, unobserved location attributes may be correlated with observed characteristics such as housing prices, potentially biasing preference estimates. Similarly, transportation costs could

[25] https://dashboard.transitmatters.org/orange/

[26] https://www.bostonglobe.com/2023/09/27/metro/new-green-line-slow-zones-pain-riders/

respond to local demand patterns, introducing another source of endogeneity. We argue that our empirical strategy and institutional context effectively address both concerns.

First, our empirical strategy for estimating location preferences sidesteps traditional endogeneity concerns. Rather than attempting to separately identify the effects of individual location attributes such as housing prices or school quality, we estimate location-specific utilities that capture the composite value of both observed and unobserved characteristics. Since we don't attempt to isolate effects of individual characteristics, the typical price endogeneity concerns do not arise in our estimation of residential location preferences.

While our approach aids in identification, it prevents us from estimating the elasticity of housing demand. As a result, we do not simulate a new market equilibrium that incorporates both consumer price responsiveness and housing supply responses. Instead, we construct bounds by considering two extreme scenarios: one in which the entire adjustment occurs through prices (with the housing stock held fixed), and another in which the adjustment occurs entirely through changes in the housing stock (with prices held constant).

Turning to transportation costs, we argue that mode-specific prices are plausibly exogenous to local residential choices. The MBTA's fare structure exhibits several features that suggest prices are not optimized in response to local demand patterns: Bus and subway fares are fixed regardless of distance traveled, while commuter rail fares follow a simplified zonal pricing scheme that is uniform across all lines. Moreover, historical zone fare adjustments appear to follow basic administrative heuristics, primarily reflecting proportional increases modified to achieve round-number fares, rather than responding to spatial variation in demand, as illustrated in Figure 1. Costs of driving—mileage and parking costs—are also plausibly exogenous with respect to work locations. Gasoline prices are determined in national markets, insulating them from Boston-specific or census-tract-specific demand fluctuations. Local parking prices largely reflect underlying land values, which are influenced by various factors beyond the direct influence of localized transportation demand.

4.3 Model Estimates

We obtain estimates of the preference parameters (β_g) for the transportation characteristics and commuting-independent home location utilities (δ_{gh}) by maximizing the likelihood in Equation (6).[27] Table 2 presents model estimates for all travel components for low-income (column 1) and high-income (column 2) households. Note that all coefficients are statistically significant (p<0.01). Several patterns become clear.

First, driving time creates greater disutility than public transit time for both income groups. High-income consumers, for example, view one minute of driving as equivalent to 1.643 minutes on public transit (-0.069/-0.042), indicating they would accept spending 64.3 percent longer on transit to avoid driving. Low-income consumers show nearly identical preferences, accepting 67.4 percent longer transit times to avoid driving. We interpret this lower public transit time disutility as stemming from passengers' ability to multitask, similar to what autonomous vehicles allow; and we leverage this in counterfactual simulations by assuming time spent in SDVs generates similar utility to public transit time.

Second, the mode-specific indicator coefficient reveals a significant fixed disutility associated with public transit use, with pronounced heterogeneity across income groups. Converting these estimates to driving-time equivalents, we find that the fixed disutility of transit usage corresponds to 0.86 minutes of driving time for lower-income households (-0.066/-0.077) and 6.29 minutes (-0.434/-0.069) for higher-income households.

Third, while the price coefficients may appear small, the resulting elasticities are similar to recent literature.[28] Importantly, price sensitivity exhibits expected heterogeneity across

[27] The gradient is described in Appendix Section A.1.

[28] The estimated price elasticities for each mode—driving, transit, combination of both—are -0.22, -0.37, and -0.55. Almagro et al. (2024) found average transit price elasticities of about -0.5 in Chicago.

income groups. While most coefficients are similar in magnitude between groups, the price coefficient for low-income households (-0.042) is about two thirds larger than for high-income households (-0.025), suggesting that the (cheaper) public transit option is more attractive for low-income workers. Together, the mode-specific indicator and the price coefficients yield mode choice patterns across income groups consistent with publicly available data: Public transit commuters tend to be more likely to have earnings below $15,000 and less likely to have salaries over $100,000 (American Public Transportation Association, 2017).

Finally, the estimates also reveal substantial transfer penalties in transit journeys (see Kreindler et al., 2023). A single transfer generates disutility equivalent to 3.5 minutes of transit time for low-income users (0.159/0.046) and 3.7 minutes for high-income users (0.157/0.042). Since the expected transfer time is already included in total trip duration, this disutility reflects other factors like schedule uncertainty, required attention, and work disruption during transfers. Our estimates carry implications for autonomous vehicles: by offering point-to-point service, SDVs eliminate transfer penalties while allowing multitasking.

4.4 Empirical Validation

The model infers preferences for transportation mode characteristics without incorporating actual transportation choices. Here, we use external survey data to assess the validity of the implied transportation preference parameters and the fit of the resulting choices, in three analyses.[29] First, we compare the model's predicted share of transit users with observed aggregate public transit usage data. Second, we compare model-predicted and survey-estimated transit ridership patterns across census tracts. Third, we compare the model's predicted use of multi-modal commuting options with survey data.

[29] This validation process presents certain challenges as our consideration of only those individuals employed in the top 100 census tracts in the Boston area by worker population creates a discrepancy between our sample and external data sources. Nevertheless, a comparative analysis may offer some assurances.

In our first validation exercise, we compare the model's predicted share of residents using public transit in their commute with data from the 2014-2018 American Community Survey (ACS) 5-year sample (U.S. Census Bureau, 2018), restricting the ACS data to the same set of residential census tracts as in our model. Among the set of home census tracts appearing in both datasets, the model predicts a 14.35 percent public transit usage rate, while the ACS data indicate an 11.1 percent rate.

Next, we compare our predicted public transit usage with survey data at the census tract level. Specifically, we create groups of 50 home location tracts each, grouped according to our model's predicted public transit usage share. For each of these groups, we compare the estimated percentage choosing public transit to the corresponding shares from the ACS. Figure 3 shows that the model predictions (x-axis) track external survey data (y-axis) relatively closely, although there are some discrepancies in census tracts with high predicted public transit usage.

Finally, we test our model with external data on station access behavior among commuter rail passengers. Our model predicts that 60.5 percent of commuter rail passengers use private vehicles to access stations. This prediction aligns well with a 2008-2009 MBTA passenger survey (Humphrey et al., 2010), which found that 65.3 percent of passengers reached stations via private motorized modes (including personal vehicles, carpools, and private shuttles).

The close alignment of the model predictions and external data on these three dimensions suggests that the model reflects reported transportation choices despite not using this information directly in the estimation.[30] This suggests that residential and workplace location data alone can be used to analyze transit usage patterns and counterfactual policies, when transit ridership data are not readily available.

[30]We repeat the validation exercises under an alternative assumption that excludes the outside transportation option, thus forcing the worker to commute using one of the three inside options. Under this assumption, the model aligns less closely with the auxiliary data, yet still yields similar counterfactual results. See Appendix Section A.2 for further details.

5 Equilibrium Effects of SDVs and Investment

Our objective is to compare the transportation outcomes resulting from SDVs with those from investments in traditional public transit. We achieve this through counterfactual simulations, modeling SDVs as having identical transportation features as traditional vehicles—except that the disutility of time spent traveling is set equal to that of passive public transit travel. Each simulation establishes a new joint equilibrium of congestion, home locations, and transportation mode choices.

To separate the roles of transportation mode choice and home location choice, we make various relocation assumptions. First, we analyze a scenario where home locations remain fixed at the status quo, isolating the impacts of transportation mode and resulting congestion changes. In addition, we use two other counterfactuals to provide intuitive bounds for the effects when relocation is permitted. In one case, workers can move but the local housing stock remains fixed at current levels. In the other case, workers can move and unlimited new housing can be constructed and priced equivalently to existing homes. The actual impact likely falls between these scenarios, depending on the elasticity of new housing supply.

The rest of this section details the counterfactual simulation methodology and results. Section 5.1 outlines the process for simulating counterfactual congestion and location equilibria. Section 5.2 presents transportation outcomes both with and without SDVs. Section 5.3 explores the effects of partial SDV adoption and parking price increases. Section 5.4 focuses on the impacts of public transit improvements in scenarios both with and without SDVs. Lastly, Section 5.5 discusses the policy implications of these findings.

5.1 Equilibrium of Home Locations and Congestion

To simulate counterfactual spatial equilibria, we must account for both residential location choices and traffic congestion effects simultaneously. Our approach uses a zonal congestion model that links vehicle miles traveled (VMT) to average traffic speeds. We solve for the equilibrium by iteratively updating zone-specific speeds until they produce home location and commuting patterns that generate VMT levels consistent with those speed assumptions. This subsection describes the congestion model and parameter calibration, as well as the equilibrium solution method.

5.1.1 Congestion Model

We model traffic congestion slowdowns using a formula proposed by Underwood (1960), which has been found to fit empirical data reasonably well (Qu et al., 2015; Wang et al., 2011). Specifically, Underwood (1960) expresses travel speed as:

$$v = v_f \exp(-k/k_0), \tag{7}$$

where v and v_f are the traffic speed with and without congestion, respectively, and k and k_0 measure road usage (VMT) and capacity, respectively.

We model traffic congestion in three concentric zones (z), each forming a ring radiating outward from Boston's City Hall. In counterfactual analyses, we find internally consistent zone-specific counterfactual VMT (k_z) and travel speeds (v_z), given calibrated values of zone specific free-flowing speed (v_{fz}) and road capacity (k_{0z}). We find these calibrated values in three steps: We first estimate the travel speeds v_z and v_{fz} in the current environment from observed travel times. Second, we infer status quo VMT (k_z) from our model. Third, given these three calibrated values, we infer corresponding road capacities k_{0z} via Equation (7).

5.1.2 Calibration

To estimate current traffic speeds, we collect travel times t_ℓ for routes ℓ between 100 randomly selected home-work location pairs from Google Maps on a weekday at two different times of day: first, to arrive at 9am (to capture v_z), and second, when leaving at 3am (to capture v_{fz}). Note that the total travel time t_ℓ for route ℓ is the sum of the travel times in each traffic zone $t_{\ell z}$; and in each zone, travel time is a function of distance and speed: $t_{\ell z} = \frac{d_{\ell z}}{v_z}$. Therefore, we can express the total travel time as:

$$t_\ell = \sum_{z' \in Z} \left(\frac{1}{v_{z'}} \right) d_{\ell z'}. \tag{8}$$

The formulation derived above lends itself naturally to a regression specification. We observe both the total trip time t_ℓ and within-zone distances $d_{\ell z}$, so we estimate the inverse of travel speeds in each zone $\left(\frac{1}{v_z} \right)$—both with and without traffic—as coefficients from regressions of t_ℓ on the zone-specific distances $d_{\ell z}$.[31]

We then use these speed parameters to infer the remaining parameters, usage k_z and road capacity k_{0z}. In the second step, we simulate the status quo VMT in each travel zone—the status quo road usage k_z—according to our model in Section 4. Finally, in the third step, we derive the zone-specific road capacity k_{z0} by solving the formula in Equation (7).

5.1.3 Equilibrium Simulation

The non-congestion values of v_{fz} and k_{0z} allow us to find the equilibrium of home locations, transportation choices, and implied traffic speeds (v_z) in counterfactual environments. To find v_z in counterfactual environments, we use an iterative process. We start with an initial

[31] We measure both distance and velocity as the components along the straight-line distance. See Appendix A.3 for details.

guess of traffic speeds, $\tilde{v}_z$, for each traffic zone. Based on these speeds, we simulate home and transportation choices, which yield a corresponding VMT k_z in each zone. Using these usage values, we apply Equation (7) to calculate updated traffic speeds, $\hat{v}_z$. These updated speeds then serve as the new initial guesses, and the process repeats until the difference between successive guesses of traffic speeds converges, yielding a consistent equilibrium of home location choices and traffic speeds.

In the status quo, the estimated speeds achieved during rush hour relative to free-flowing traffic ($^{v}/_{v_f}$) are 58 percent in the inner ring (closest to Boston center), 65 percent in the middle ring, and 97 percent in the outer ring. In our counterfactuals, below, we show that SDVs cause VMT to increase, resulting in more pronounced slowdowns.

5.2 Primary Impacts of Self-Driving Vehicles

In our first counterfactual analyses, we examine how transportation mode utilization and consumer welfare change when SDVs are introduced and fully adopted, therefore reducing the time disutility of driving to that of public transit. Table 3 summarizes the simulation results. The first column reflects the baseline scenario under current conditions, while the subsequent columns summarize counterfactual simulations incorporating SDVs. Columns 2 through 4 endogenize traffic speeds to respond to changing congestion levels. The second column isolates the effects of SDVs on transportation mode choice by preserving status-quo home locations.[32] In the third column, relocation is permitted but the income-type-specific housing supply at the census tract level is fixed at status quo levels;[33] and the fourth column

[32] To preserve home location choices, we fix the predicted probabilities of workers selecting each home location to align with the status quo. We achieve this by jointly determining equilibrium congestion levels and type-home-work-specific inclusive values ($\tilde{\delta}_{ghw}$) to match the baseline probabilities observed without SDVs.

[33] The procedure of adjusting type- and home-specific inclusive values ($\tilde{\delta}_{gh}$) resembles that in footnote 32, except that the inclusive values are now fixed at the group- and home-specific status quo level. We model separate housing markets for low- and high-income households to reflect real-world policies such

allows for new construction at prices consistent with existing housing. These latter two scenarios establish lower and upper bounds for the equilibrium effects of autonomous vehicles, where the exact impacts depend on housing supply elasticities. Finally, the fifth column incorporates housing supply adjustments—similar to Column 4—but maintains traffic speeds at their baseline levels, reflecting a scenario where autonomous vehicles achieve more efficient road space utilization (Hyldmar et al., 2019).

We first consider the impacts of SDVs on miles driven and home location choices—shown in the first two rows of Table 3—focusing on counterfactual scenarios which adjust congestion. Across all three scenarios, average miles driven rise substantially. The increase in miles driven ranges from 26 percent (9.54 to 11.98 miles) when home locations are fixed, to 41 percent (9.53 to 13.44 miles) when the housing supply is perfectly flexible. When home locations are fixed, the mileage changes arise entirely from changes in commuting mode choices. Figure 4 illustrates these changes, separately for workers living various distance ranges from their work locations. Both driving and multi-modal transportation shares rise at the expense of pure public transit commuting and hybrid work. Moreover, while these patterns hold for all distance bands, they are particularly pronounced the further a person lives from their workplace.[34]

The effect of SDVs on miles driven rises further when we allow workers to move and housing supply to adjust, as workers move further from their work locations when driving is made less burdensome: with fully flexible housing, the average distance from work rises by 16 percent. We illustrate these relocations in Figures 5 and 6. In Figure 5, we depict the changes in housing prices necessary to keep the housing supply unchanged (corresponding to Column 3 of Table 3), and Figure 6 shows changes in residential density if housing prices

as Massachusetts legislation that mandates municipalities to maintain a share of affordable housing units restricted to residents earning below 80 percent of the county's median income, and the tendency for different income groups to select homes of different quality and size.

[34]The figure resembles the analogous figure when instead assuming elastic housing supply.

remained constant (Column 4). Both figures reflect large increases in housing appeal in the outer suburbs of Boston relative to areas closer to the city center. This finding suggests that SDVs could lead to further urban sprawl, similar to the effect of car adoption over the past century (Redding, 2022; Glaeser and Kahn, 2004), while also potentially reducing the upward pressure on housing prices near city centers observed in recent decades (Glaeser et al., 2012; Couture et al., 2024).

The next three rows of Table 3 show the impacts of SDVs on the public transit system. When housing locations are persistent, total public transit revenue per worker decreases by 19 percent (from \$0.79 to \$0.64 per day). This decrease arises from both a decrease in the revenue per rider (from \$5.51 to \$4.88) and a decrease in the share of workers who use public transit (from 14.4 percent to 13.1 percent). Allowing home locations to adjust mitigates the decrease in revenue: When the housing supply is fully flexible—leading to longer distances to work, more congestion, and more workers driving to transit access points—public transit revenue decreases by only 13 percent. The smaller revenue reduction occurs because the ridership decline is less severe and the revenue per rider reverts towards status quo levels.

The above discussion provides an equilibrium analysis in the sense that it accounts for changes in traffic speeds resulting from increased road utilization; and the bottom panel of Table 3 shows significant speed decreases in Columns 2 through 4 relative to the status quo, especially in the inner and middle rings. There is, however, a possibility that the new technologies will also allow for more efficient traffic flows. To examine this possibility, Column 5 of Table 3 examines a scenario in which individuals can relocate freely—as in Column 4—but traffic speeds are held constant at baseline levels. This scenario suggests an even greater shift toward driving (raising miles driven by 72 percent relative to the status quo, and by 22 percent relative to Column 4), as well as home choices even farther from work locations. These changes also imply a sharper decline in public transportation use, with the share of users dropping to 12.9 percent.

Our model also allows for welfare evaluation under the assumption that inherent home location utilities remain unchanged—in Columns 4 and 5 of the table. The improved transit options lead to sizable welfare gains for consumers. On average, we estimate that consumers gain \$12.23 per workday when accounting for congestion, and up to \$21.39 when keeping traffic speeds unchanged. The gains are largest for high-income consumers, whose consumer surplus rises by \$15.71 (or \$27.30 without congestion), whereas low-income consumer surplus rises by just \$4.95 (\$9.04). The welfare effects reflect the benefits of a hands-free commute and the cost savings when households relocate to more distant areas with more affordable housing. These results suggest that addressing congestion challenges could enhance the commuter benefits of SDVs.

5.3 Robustness to Changing Environments

The above analysis relies on two key assumptions regarding the adoption of SDVs: that SDVs completely replace conventional cars, and that increased car usage does not affect other costs associated with driving—such as parking costs. In this section, we explore how these assumptions affect our results.

5.3.1 Rising Shares of SDVs

Adoption of SDVs may be slow for various reasons. Here, we examine the effects from above when SDVs account for certain shares of the full fleet of cars, ranging from zero to 100 percent, using the specification where housing supply is flexible.[35] The left panel of Figure 7 shows how average miles driven and transit revenue change as we raise the share of SDVs.

[35]We assume a certain fraction of consumers are endowed with SDVs, while remaining consumers lack access.

The relationships between the share of SDVs and both outcomes are almost linear, with values bounded between the corresponding values in Columns 1 and 4 of Table 3. The same patterns also hold for the other outcomes reported above. Moving forward, we continue to assume full adoption of SDVs, noting that the results are roughly proportional for lower take-up rates.

5.3.2 Rising Parking Costs

The increasing use of cars for commuting can lead to various equilibrium shifts, most notably a higher demand for workplace parking, which may drive up parking prices. To explore this, we analyze changes in miles driven and transit revenue across a range of parking price scenarios, from current levels (our baseline) to a doubling of prices. As shown in Panel B of Figure 7, even a doubling of parking costs only has a modest impact on driving behavior: miles driven would rise by approximately 33 percent, compared to just over 40 percent at baseline prices. In contrast, transit revenue is more responsive to parking costs. At current parking prices, we project revenue to decline by about 12 percent. However, if parking prices increase by roughly 80 percent, transit revenue would stabilize at pre-SDV levels, and would begin to rise with further increases.

5.4 Comparative Analysis of Transit Improvements

The apparent substitution from transit to self-driving vehicles raises an important question for policy makers and public investors: Will improvements to public transit still yield meaningful changes to transit usage, road travel, and revenue? To address this question, we examine the effects of public transit improvements on revenue and VMT, both in a world with and without SDVs. Here, we model improvements in train speeds as reductions in onboard (seated) travel time, while holding fixed both the waiting times associated with transit and the disutility arising from transfers at status quo levels.[36] Table 4 reports counterfactual percentage changes in miles driven (Columns 1 and 2) as well as in transit revenue per resident (Columns 3 and 4). All numbers are percent changes relative to the baseline without SDVs or transit improvements (Column 1 of Table 3).

Consider first the impacts on miles driven. The first column shows effects when SDVs remain unavailable. In this scenario, consumers drive less as public transit trains become faster. For example, if the time spent on public transit were to decrease by 40 percent, some consumers would switch to public transit and miles driven would decrease by 3.36 percent (se=0.12). We find somewhat similar marginal impacts in a scenario with SDVs (Column 2). For example, improving transit speeds by 40 percent reduces VMT by about 2.3 percent $\left(\approx \frac{140.90-137.68}{140.90} \times 100\right)$, suggesting that investment can continue to yield behavioral changes.

However, the decreases in miles driven pale in comparison with the impact of SDV adoption. Even if the time on the train were eliminated completely—de facto allowing for teleportation between entry and exit stations—miles driven would still be 26.71 percent (se=1.87) higher than under the status quo without transit improvements or SDVs. These findings reflect the persistence of mode choice, combined with large increases in mileage from moving further from one's place of work.[37]

This qualitative result—that mileage reductions from public transit improvements are still smaller than mileage increases due to SDVs—appears to hold even with relatively low SDV adoption. Because the effects of SDVs are relatively linear in their adoption rate, we can use back of the envelope calculations to infer a threshold penetration rate ϕ of SDVs

[36] In Appendix Table A.4, we present improvements in train station quality, convenience, and certainty associated with transfers (which we model as reductions in the hassle costs and time spent waiting during transfers). All results remain, although traffic stop improvements, which account for a small share of the total commuting time, have slightly smaller effects.

[37] This fairly strong persistence arises from a relatively parsimonious model. If we also introduced a random coefficient on consumer preference for a specific mode of transportation (e.g., transit), this persistence would likely become even stronger.

beyond which VMT increase for a given level of speed improvement. We find this threshold by solving for weights that make a weighted sum of the Column-1 effect and the Column-2 effect equal to zero. For example, the aggregate joint impacts of SDVs and a 40 percent reduction in time between entry and exit stations would still be positive—more mileage—if SDV penetration exceeded 8.2 percent.[38] Even with 100 percent speed improvements, implying instant travel along a given transit segment, VMT would still exceed baseline levels for any penetration rate above 34.9 percent.

While the first two columns of Table 4 indicate that improvements in public transportation cannot fully reverse the increase in VMT from SDVs, Columns 3 and 4 suggest different outcomes for revenues. For example, a 20 percent improvement in train speed would raise transit revenue by 27.7 percent without SDVs; and even with SDVs, revenues are 10.3 percent above baseline levels. Similar patterns ensue from reducing the hassle cost of transfers.

Finally, while SDVs influence the overall revenues, the rate at which revenues respond to transit improvements remains largely unaffected. For example, in the presence of SDVs, improving transit speeds by 20 percent raises revenues by 26.6 percent, only slightly below the effect of the same improvement in the absence of SDVs.[39] This consistency underscores the continued relevance of public transit improvements as a policy tool for managing urban mobility.

5.5 Policy Considerations

The effectiveness of transit improvements in the presence of SDVs depends on the policy-makers' objectives. If the goal is to recover the fixed costs of these enhancements through fare revenues, then SDVs have a relatively small impact on the viability of such transit im-

[38] With a 40 percent reduction in passive transit travel time, the overall change in mileage with SDV penetration rate ϕ equals $(1-\phi) \times -3.36 + \phi \times 37.68$. This expression exceeds zero when $\phi \geq 0.0819$.

[39] $\frac{10.25\%+12.91\%}{100\%-12.91\%} \approx 26.6\%$.

provements. However, our findings are less encouraging for policymakers seeking to mitigate increases in private vehicle miles traveled, whether motivated by environmental externalities, public health concerns, congestion externalities, or other considerations. In this case, one potential policy response is transit subsidization, potentially extending to fare-free service, though this approach would require a reconsideration of how public transit is funded. Another approach is to promote generous work-from-home policies and enhanced incentives for carpooling.

If the proverbial "carrot" approaches of increasing the appeal of non-driving options fail to limit car use, policymakers may need to consider regulatory "sticks." Prior research documents the efficacy of various driving regulations, including congestion pricing and road access restrictions (see Section 2.4). The similarity of the impacts of investments in public transit in scenarios with and without SDVs in our study suggests that other interventions may also have similar effects regardless of SDV adoption, indicating that these policy tools may remain effective even in a future transportation ecosystem that includes self-driving vehicles.

Ultimately, a mix of policies may be necessary to effectively shape urban sprawl and transportation in an SDV-driven world. Strategic investments in transit, regulatory interventions, and incentives to reduce driving are complementary tools to balance the consumer benefits of SDVs against their broader societal costs.

6 Conclusion

New and emerging driving technologies may change the role of public transit looking forward. Self-driving vehicles (SDVs) are already available as ride-hailing services in several metropolitan areas, and further advances in the availability of autonomous driving technologies suggest that it may be only a matter of time until SDVs are widely available to the public. Still, local and federal governments continue to invest heavily in public transit. This raises three main questions: Will SDVs substantially change the way that we commute; will investments in public transit remain effective; and how do these relative effects compare?

This paper provides a first answer to these questions. In the context of the Boston region, we find, first, that the widespread adoption of SDVs significantly changes both home locations and transportation choices: workers will likely move further away from their workplace and become more likely to use cars for commuting, leading to increased congestion and lower public transit revenue. Second, we find that improvements in the public transit experience continue to meaningfully affect consumer behavior. Third, while these improvements cannot reverse the increase in miles driven effected by SDVs, their positive effects on public transit revenue *can* outweight those of SDVs if the improvements are large enough. These results suggest that public transit can remain an important and valuable means of transportation looking forward.

Our paper captures just one way in which SDVs will change daily live. We focus on commuting to work, but SDVs may also take over other parts of life, such as deliveries, visits to friends and family, and nights out. Dealing with increased demands for cars, and limiting their effects on the environment, could become an important consideration for public policy.

References

Agrawal, A., J. S. Gans, and A. Goldfarb (2019): "Artificial intelligence: the ambiguous labor market impact of automating prediction," Journal of Economic Perspectives, 33, 31–50.

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

Almagro, M., F. Barbieri, J. C. Castillo, N. G. Hickok, and T. Salz (2024): "Optimal Urban Transportation Policy: Evidence from Chicago," Tech. rep., National Bureau of Economic Research.

American Public Transportation Association (2017): "Who Rides Public Transportation," .

Barwick, P. J., S. Li, A. Waxman, J. Wu, and T. Xia (2024): "Efficiency and equity impacts of urban transportation policies with equilibrium sorting," American Economic Review, 114, 3161–3205.

Baum-Snow, N. (2007): "Did highways cause suburbanization?" The quarterly journal of economics, 122, 775–805.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," Econometrica: Journal of the Econometric Society, 841–890.

Brancaccio, G., M. Kalouptsidi, T. Papageorgiou, and N. Rosaia (2023): "Search frictions and efficiency in decentralized transport markets," The Quarterly Journal of Economics, 138, 2451–2503.

Brooks, L. and Z. Liscow (2023): "Infrastructure costs," American Economic Journal: Applied Economics, 15, 1–30.

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

Buchholz, N., L. Doval, J. Kastl, F. Matejka, and T. Salz (2025): "Personalized Pricing and the Value of Time: Evidence from Auctioned Cab Rides," Econometrica, 93, 929–958.

Chen, M. K., P. E. Rossi, J. A. Chevalier, and E. Oehlsen (2019): "The value of flexible work: Evidence from Uber drivers," Journal of political economy, 127, 2735–2794.

Conlon, C. and J. Gortmaker (2020): "Best Practices for Differentiated Products Demand Estimation with PyBLP," Working Paper.

Couture, V., G. Duranton, and M. A. Turner (2018): "Speed," Review of Economics and Statistics, 100, 725–739.

Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2024): "Income growth and the distributional effects of urban spatial sorting," Review of Economic Studies, 91, 858–898.

Cramer, J. and A. B. Krueger (2016): "Disruptive change in the taxi business: The case of Uber," American Economic Review, 106, 177–182.

Davis, M. A., A. C. Ghent, and J. Gregory (2024): "The work-from-home technology boon and its consequences," Review of Economic Studies, 91, 3362–3401.

Delventhal, M. and A. Parkhomenko (2024): "Spatial implications of telecommuting," Available at SSRN 3746555.

Duranton, G. and M. A. Turner (2011): "The fundamental law of road congestion: Evidence from US cities," American Economic Review, 101, 2616–2652.

Durrmeyer, I. and N. Martinez (2024): "The welfare consequences of urban traffic regulations," Tech. rep., TSE Working Paper.

Fajgelbaum, P. D. and E. Schaal (2020): "Optimal transport networks in spatial equilibrium," Econometrica, 88, 1411–1452.

Frechette, G. R., A. Lizzeri, and T. Salz (2019): "Frictions in a competitive, regulated market: Evidence from taxis," American Economic Review, 109, 2954–2992.

Glaeser, E. L., J. D. Gottlieb, and K. Tobio (2012): "Housing booms and city centers," American Economic Review, 102, 127–133.

Glaeser, E. L. and M. E. Kahn (2004): "Sprawl and urban growth," in Handbook of regional and urban economics, Elsevier, vol. 4, 2481–2527.

——— (2010): "The greenness of cities: Carbon dioxide emissions and urban development," Journal of urban economics, 67, 404–418.

Glaeser, E. L., M. E. Kahn, and J. Rappaport (2008): "Why do the poor live in cities? The role of public transportation," Journal of urban Economics, 63, 1–24.

Graham, M., E. McEntarfer, K. McKinney, S. Tibbets, and L. Tucker (2022): "LEHD Snapshot Documentation, Release S2021_R2022Q4," Tech. rep.

Guda, H. and U. Subramanian (2019): "Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives," Management Science, 65, 1995–2014.

Hall, J. D., C. Palsson, and J. Price (2018): "Is Uber a substitute or complement for public transit?" Journal of urban economics, 108, 36–50.

Hsu, W.-T. and H. Zhang (2014): "The fundamental law of highway congestion revisited: Evidence from national expressways in Japan," Journal of urban economics, 81, 65–76.

Humphrey, T. J., S. P. Andrews, M. A. Maslova, and M. McShane (2010): "MBTA Systemwide Passenger Survey: Commuter Rail 2008–09," Technical report, Boston Region Metropolitan Planning Organization, Boston, MA.

Hyldmar, N., Y. He, and A. Prorok (2019): "A fleet of miniature cars for experiments in cooperative driving," in 2019 International Conference on Robotics and Automation (ICRA), IEEE, 3238–3244.

Ioannides, Y. M., H. G. Overman, E. Rossi-Hansberg, and K. Schmidheiny (2014): "The effect of information and communication technologies on urban structure," Economic Policy, 23, 202–242.

Jerch, R., P. J. Barwick, S. Li, and J. Wu (2024): "The impact of road rationing on housing demand and sorting," Journal of Urban Economics, 140, 103642.

Kahn, M. E. and R. Walsh (2015): "Cities and the Environment," Handbook of regional and urban economics, 5, 405–465.

Kreindler, G. (2024): "Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications," Econometrica, 92, 1233–1268.

Kreindler, G., A. Gaduh, T. Graff, R. Hanna, and B. A. Olken (2023): "Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta," Tech. rep., National Bureau of Economic Research.

Larson, W. and W. Zhao (2020): "Self-driving cars and the city: Effects on sprawl, energy consumption, and housing affordability," Regional Science and Urban Economics, 81, 103484.

Monte, F., C. Porcher, and E. Rossi-Hansberg (2023): "Remote Work and City Structure," .

Neil, D. (2023): "Rumble Seat: Mercedes-Benz Introduces Hands-Free Driving. How Does It Compare to Tesla?" The Wall Street Journal, uRL: https://www.wsj.com/lifestyle/cars/mercedes-benz-hands-free-drive-pilot-review-tesla-comparison-ede47ac7, Accessed: 2023-11-16.

Qu, X., S. Wang, and J. Zhang (2015): "On the fundamental diagram for freeway traffic: A novel calibration approach for single-regime models," Transportation Research Part B: Methodological, 73, 91–102.

Redding, S. J. (2022): "Suburbanization in the USA, 1970–2010," Economica, 89, S110–S136.

Severen, C. (2023): "Commuting, labor, and housing market effects of mass transportation: Welfare and identification," Review of Economics and Statistics, 105, 1073–1091.

Smith, A. and M. Anderson (2017): "Americans' attitudes toward driverless vehicles," Automation in Everyday Life.

Tarduno, M. (2021): "The congestion costs of Uber and Lyft," Journal of Urban Economics, 122, 103318.

Tsivanidis, N. (2023): "Evaluating the impact of urban transit infrastructure: Evidence from Bogota's Transmilenio," University of California, Berkeley.

Underwood, R. T. (1960): "Speed, volume, and density relationships," .

U.S. Census Bureau (2018): "American Community Survey (ACS) 5-Year Estimates, 2014-2018," https://www.census.gov/programs-surveys/acs/, data accessed via U.S. Census Bureau API.

Van Ommeren, J. and P. Rietveld (2005): "The commuting time paradox," Journal of Urban Economics, 58, 437–454.

Wang, H., J. Li, Q.-Y. Chen, and D. Ni (2011): "Logistic modeling of the equilibrium speed–density relationship," Transportation research part A: policy and practice, 45, 554–566.

Weisbrod, G. and A. Reno (2009): Economic impact of public transportation investment, American public transportation association Washington, DC.

7 Tables and Figures

Table 1: Summary Statistics, Weighted by Home/Work Location Frequency

	N	Mean	Std. Dev.	Median	Min	Max
Public Transit Only						
Duration (minutes)	52,132	74.41	42.83	66.00	0.00	281.00
# segments	52,132	2.18	1.23	2.00	0.00	8.00
Price ($)	52,132	4.37	4.57	1.70	0.00	30.20
Driving to Transit						
Duration (transit)	97,735	50.35	26.96	46.00	0.00	168.00
Duration (driving)	97,735	35.20	17.90	31.69	0.00	238.16
# transit segments	97,735	1.69	0.93	2.00	0.00	6.00
Price ($)	97,735	16.66	6.14	16.31	1.00	46.65
Driving						
Duration	105,770	28.86	21.03	23.61	0.00	236.14
Price ($)	105,770	12.60	7.66	12.92	0.00	41.68
Fastest Commute Mode						
Duration	105,770	28.86	21.03	23.60	0.00	236.14

Notes: This table shows summary statistics for varying transportation mode options, for consumers living and working in census tracts near Boston, MA. All variables are weighted by the number of workers observed in each home/work location pair. "Public Transit Only" statistics are calculated from all origin-destination pairs where commuting exclusively by public transit is feasible; "Driving to Transit" includes the subset for which a combination of driving and transit is possible; and "Driving" describes all pairs for which driving is possible.

Table 2: Parameter Estimates

	Households Earning:	
	$\leq$ 3,333 per month	$>$ \$3,333 per month
Constant	5.090	5.452
	(0.024)	(0.026)
1(Involves Public Transit)	-0.066	-0.434
	(0.028)	(0.019)
Transfer Count	-0.159	-0.157
	(0.024)	(0.017)
Driving Duration	-0.077	-0.069
	(0.000)	(0.000)
Transit Duration (Non-Auto)	-0.046	-0.042
	(0.001)	(0.001)
Price (\$)	-0.042	-0.025
	(0.001)	(0.001)

Notes: This table reports model estimates of the preference parameters for location and transportation characteristics from the demand model described in Section 4. The model is estimated separately for low-income households (left column, with incomes at \$3,333 or below) and high-income households (right column, with monthly incomes above \$3,333). Standard errors are reported in parentheses. All parameters are statistically significant at the 1 percent level.

Table 3: Simulated Outcomes of Self-Driving Vehicles (SDVs) Adoption

		Self-Driving Vehicles (SDVs)			
		Dynamic Congestion Effects			Congestion from Status Quo
	Status Quo	Persistent Location Choices	Fixed Housing Supply	Flexible Housing Supply	
E[Miles Driven (one way)]	9.54	11.98	12.70	13.44	16.42
	(0.13)	(0.12)	(0.14)	(0.23)	(0.25)
E[Distance From Work]	14.41	14.41	15.21	16.74	18.92
	(0.09)	(0.09)	(0.10)	(0.19)	(0.21)
E[Public Transit Revenue ($) Per Worker/Day]	0.79	0.64	0.65	0.69	0.64
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
E[Public Transit Revenue ($) Per Rider/Day]	5.51	4.88	5.03	5.18	4.92
	(0.04)	(0.02)	(0.03)	(0.04)	(0.03)
Percent Taking Public Transit	14.35	13.13	13.01	13.29	12.93
	(0.33)	(0.28)	(0.28)	(0.29)	(0.27)
$\frac{\textbf{Speed with Congestion}}{\textbf{Free-Flow Speed}}$					
Inner Ring	0.58	0.52	0.49	0.47	0.58
	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)
Middle Ring	0.65	0.58	0.57	0.51	0.65
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Outer Ring	0.97	0.92	0.92	0.91	0.97
	(0.05)	(0.08)	(0.08)	(0.08)	(0.05)

Notes: This table shows estimated driving, location, and public transit usage outcomes in the baseline case without SDVs (Column 1) as well as in counterfactual scenarios when cars are autonomous. Columns 2 through 4 report outcomes when changes in miles driven affect driving speeds. These columns represent different assumptions about the housing supply: Column 2 maintains each individual's original probability of residing in any given census tract, Column 3 allows residential relocation but keeps the total housing supply per census tract fixed, and Column 4 assumes perfectly elastic housing supply. Finally, Column 5 also assumes perfectly elastic housing supply but keeps traffic speeds at the status quo level. All scenarios include an outside transportation option that involves neither driving nor public transit. For robustness, Table A.2 in the appendix presents analogous outcomes when the model is re-estimated without this outside option. Standard errors in parentheses are calculated by parametric bootstrap, accounting for uncertainty in both consumer preference parameters and congestion model parameters.

Table 4: Impacts of Public Transit Speed Improvements

	Miles Driven		Transit Revenue	
	No SDVs	SDVs	No SDVs	SDVs
Status Quo Transit	0.00%	40.90%	0.00%	-12.91%
	(0.00)	(1.81)	(0.00)	(0.68)
20% Onboard Travel Time Reduction	-1.50%	39.55%	27.70%	10.25%
	(0.08)	(1.81)	(0.34)	(0.99)
40% Onboard Travel Time Reduction	-3.36%	37.68%	65.80%	42.21%
	(0.12)	(1.81)	(0.95)	(1.55)
60% Onboard Travel Time Reduction	-6.01%	35.17%	119.33%	86.87%
	(0.20)	(1.79)	(2.02)	(2.55)
80% Onboard Travel Time Reduction	-9.52%	31.64%	195.15%	150.46%
	(0.30)	(1.82)	(3.94)	(4.34)
100% Onboard Travel Time Reduction	-14.27%	26.71%	303.90%	242.22%
	(0.42)	(1.87)	(7.25)	(7.40)

Notes: This table reports simulated percentage changes in miles driven (left two columns) and transit revenue (right two columns) for varying percent improvements in the speed of transit trains, relative to status quo transit speeds with no SDVs. Improvements are applied only to onboard (seated) travel time, while waiting times and the disutility associated with transfers remain fixed at baseline levels. Columns 1 and 3 report outcomes when there are no SDVs, and Columns 2 and 4 show outcomes, relative to the status quo, in the presence of SDVs. Standard errors in parentheses are calculated by parametric bootstrap, accounting for uncertainty in both consumer preference parameters and congestion model parameters.

Figure 1: Commuter Rail Zone Prices: 1989 vs 2024

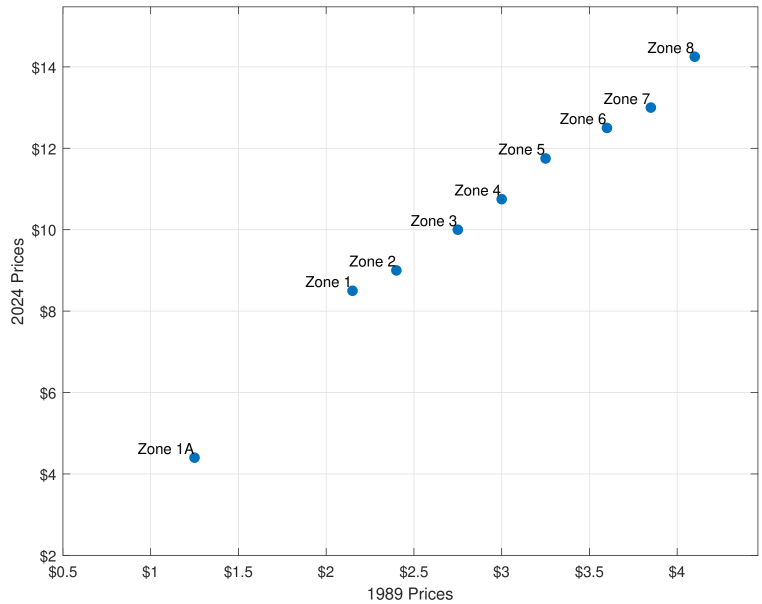

Notes: This figure compares nominal commuter rail ticket prices for commuter rail zones as defined by the Massachusetts Bay Transportation Authority (MBTA), in 1989 (x-axis) vs. in 2024 (y-axis). 1989 prices are obtained from the Commonwealth Magazine (2016). See https://commonwealthbeacon.org/uncategorized/pulling-the-data-together/.

Figure 2: Population Density by Census Tract

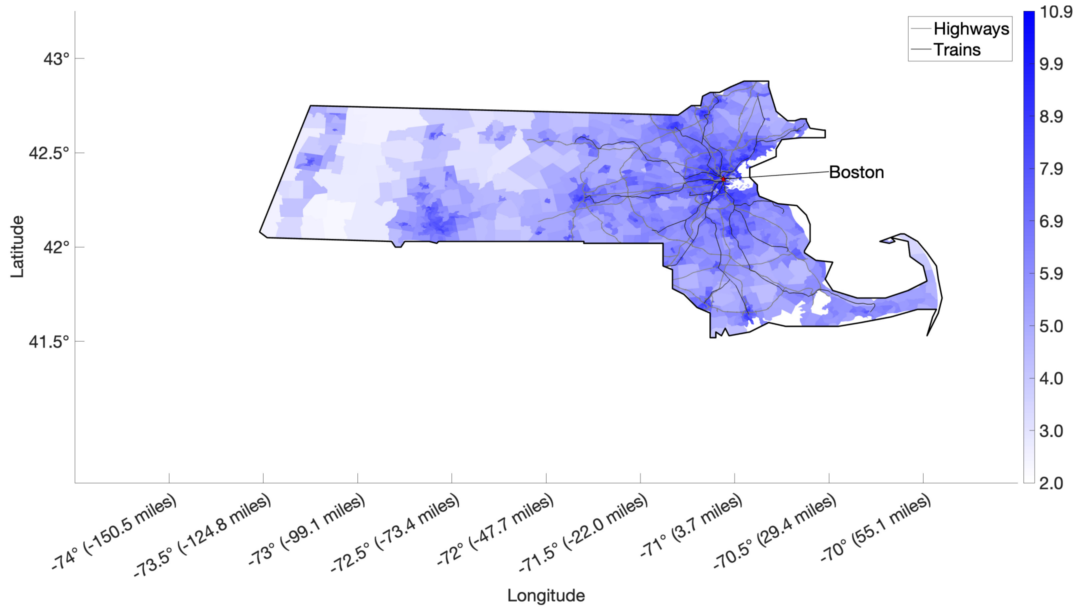

Notes: The map shows the natural log of the home-location population density in workers per square mile of land (excluding bodies of water), by census tract in Massachusetts. Darker shades of blue correspond to denser census tracts.

Figure 3: Comparison of Model Predictions and ACS Survey Data

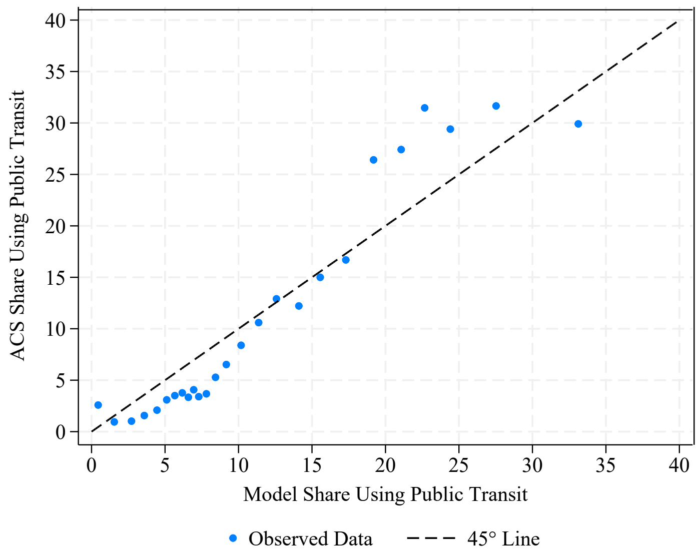

Notes: This figure compares the model's predicted census tract-level share of workers using public transit (x-axis) with 5-year estimates based on survey data from the 2014-2018 American Community Survey (ACS) (U.S. Census Bureau, 2018). Points represent bins of 50 census tracts each, grouped by predicted public transit usage percentage. The 45-degree dashed line represents perfect matches between the model's prediction and the American Community Survey's estimate.

Figure 4: Changes in Transportation Mode Choices from SDVs – Home Locations Fixed

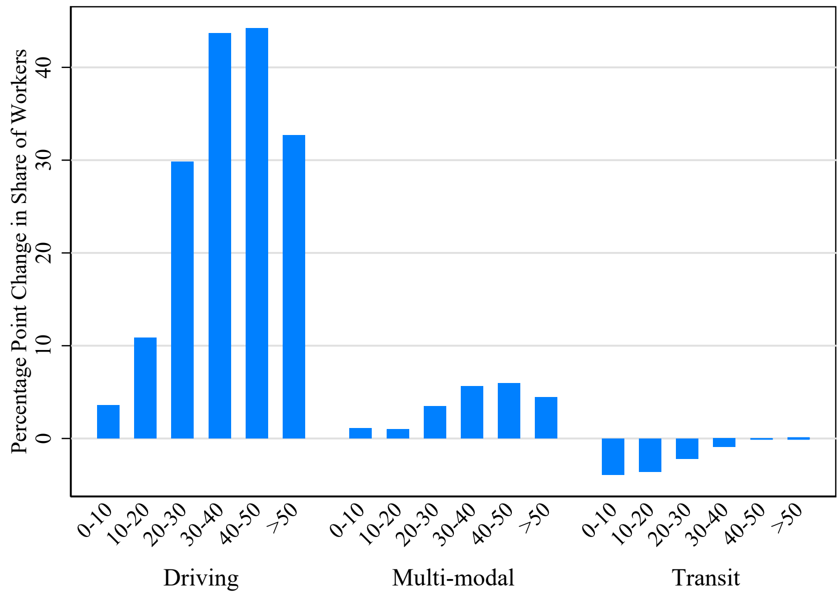

Notes: The graph shows changes in the shares of workers using each of three modes of transportation when SDVs are introduced and we retain each individual's original probability of residing in any given census tract. We report changes separately for various distance ranges, from 0 to 10 miles, to over 50 miles. For example, among workers living between 20 and 30 miles from their workplace, the share driving to work rises by 34 percentage points, the share using both cars and transit rises by about 4 percentage points, and the share using only public transit decreases by about 3 percentage points.

Figure 5: Changes in Housing Prices from SDVs with Inelastic Housing Supply

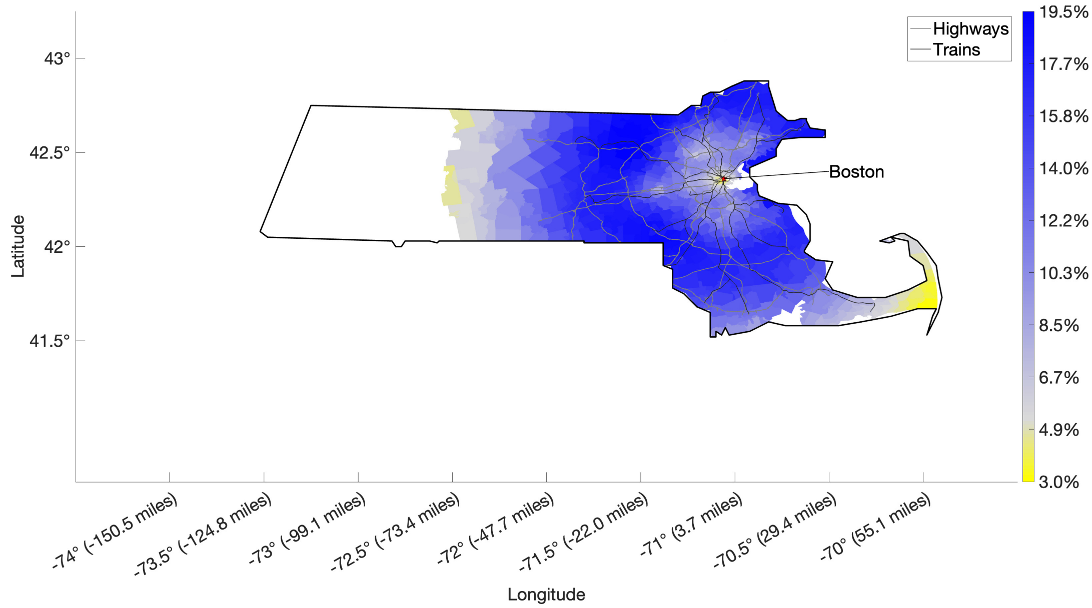

Notes: The map displays, for each census tract, the percentage change in housing prices required to balance demand with the existing housing supply when SDVs are introduced (as in Column 3 of Table 3). Blue shades indicate larger price increases, while yellow shades represent smaller price changes. Housing prices increase throughout the Boston metro area as SDVs enhance the viability of commuting, making locations within driving distance to Boston more attractive compared to remote work from more distant areas; areas moderately far from the city center increase the most. Note that the model is executed independently for high- and low-income populations. The color-coded map displayed represents predictions specifically for the low-income group. The analogous visualization for high-income individuals shows very similar patterns.

Figure 6: Changes in Worker Home Locations from SDVs with Elastic Housing Supply

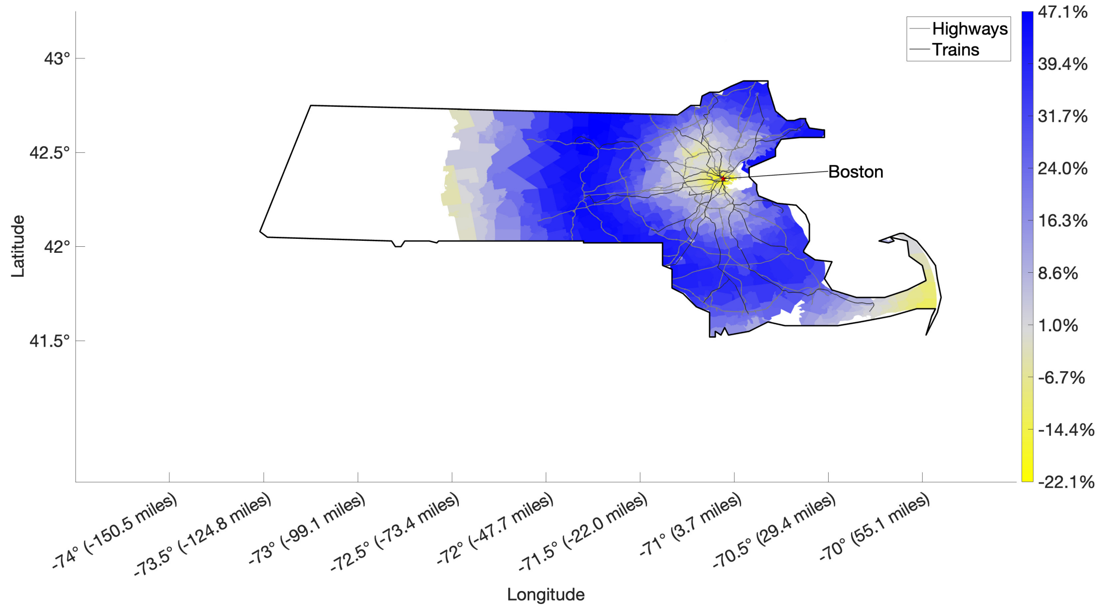

Notes: The map shows, for each census tract, the percentage change in (working) residents with the advent of SDVs when assuming perfectly elastic housing supply (as in Column 4 of Table 3). Blue hues represent increases in population, whereas yellow hues indicate reductions. Light gray hues indicate close to no change.

Figure 7: Partial Self-Driving Penetration

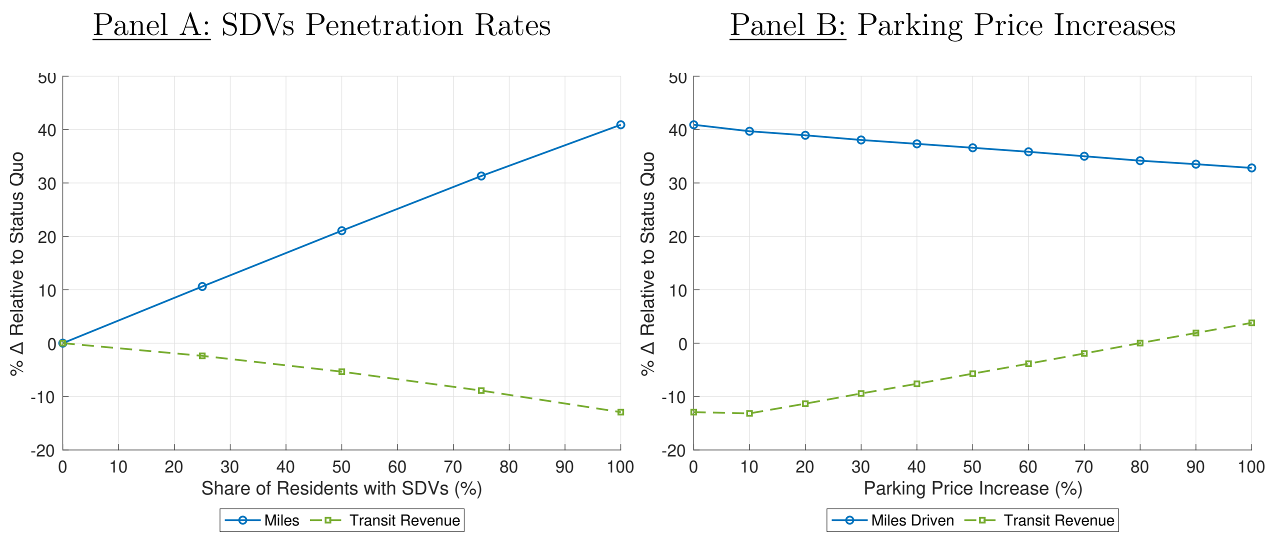

Notes: This figure shows estimated percent changes—relative to the status quo—in miles driven (solid, blue lines) and transit revenues (dashed, green lines) for varying assumptions about SDV adoption rates (Panel A) and parking prices (Panel B). All estimates assume perfectly elastic housing supply. In Panel A, the leftmost points depict the status quo with no SDVs, whereas the rightmost points describe full SDV adoption reflecting Column 4 of Table 3; and intermediate points reflect partial SDV adoption. In Panel B, the leftmost points coincide with the rightmost points in Panel A, reflecting the effects from full adoption with unchanged prices. As we move to the right in Panel B, parking prices increase and SDV adoption remains at 100 percent.

A Appendix

A.1 Gradient

The log-likelihood gradient with respect to each parameter β_{gl} can be expressed as:

$$\frac{\partial \ln(L(\beta_g))}{\partial \beta_{gl}} = \sum_{h \in H, w \in W} \#_{ghw} \frac{\left(\frac{\partial \mathbb{s}_{ghw}}{\partial \beta_{gl}}\right)}{\mathbb{s}_{ghw}}. \tag{9}$$

The derivative of market shares with respect to parameters involves both direct effects and adjustments to δ that ensure the predicted and observed market shares continue to match as the parameters β_g change:

$$\frac{\partial \mathbb{s}_{ghw}}{\partial \beta_{g\ell}} = \mathbb{s}_{ghw} \times \left(\frac{\partial \lambda(X_{hw}, \beta_g)}{\partial \beta_{g\ell}} + \frac{\partial \delta_{gh}}{\partial \beta_{g\ell}} - \sum_{h' \in H} \mathbb{s}_{gh'w} \times \left(\frac{\partial \lambda(X_{h'w}, \beta_g)}{\partial \beta_{g\ell}} + \frac{\partial \delta_{gh'}(X_{h'w}, \beta_g)}{\partial \beta_{g\ell}} \right) \right). \tag{10}$$

The direct effect has a logistic form:

$$\frac{\partial \lambda(X_{hw}, \beta_g)}{\partial \beta_{g\ell}} = \frac{\sum_{c \in C} X_{c\ell\kappa} \exp(X_{chw}\beta_g)}{1 + \sum_{c \in C} \exp(X_{chw}\beta_g)}. \tag{11}$$

Following Conlon and Gortmaker (2020), the equilibrium adjustment of home location quality levels—the indirect effect—can be found using the implicit function theorem:

$$\underbrace{\frac{\partial \delta_g}{\partial \beta_g}}_{\text{H by L}} = - \left(\underbrace{\frac{\partial \mathbb{s}_g}{\partial \delta_g}}_{\text{H by H}} \right)^{-1} \left(\underbrace{\frac{\partial \mathbb{s}_g}{\partial \beta_g}}_{\text{H by L}} \Big|_{\delta \text{ constant}} \right). \tag{12}$$

The Jacobian matrix elements have analytical forms. For the share response to home location quality:

$$\frac{\partial \mathbb{s}_{gh}}{\partial \delta_{gh}} = \int_w \mathbb{s}_{ghw} \times (1 - \mathbb{s}_{ghw}), \tag{13}$$

$$\frac{\partial \text{\$}_{gh}}{\partial \delta_{g,h' \neq h}} = \int_w -\text{\$}_{ghw} \times \text{\$}_{gh'w}, \tag{14}$$

and for the direct parameter effects holding home location quality fixed:

$$\left.\frac{\partial \text{\$}_{ghw}}{\partial \beta_{g\ell}}\right|_{\delta \text{ constant}} = \text{\$}_{ghw} \times \left(\frac{\partial \lambda(X_{hw}, \beta_g)}{\partial \beta_{g\ell}} - \sum_{h' \in H} \text{\$}_{gh'w} \times \left(\frac{\partial \lambda(X_{h'w}, \beta_g)}{\partial \beta_{g\ell}} \right) \right). \tag{15}$$

A.2 Alternative Model Specification: No Outside Option

In this section, we examine the robustness of our main results by considering an alternative specification that excludes the outside transportation option incorporated in our primary model. This outside option encompasses remote work, hybrid arrangements, secondary residences near workplaces, and similar alternatives to traditional commuting.

Table A.1 presents the estimated coefficients from the location and transporation choice model when this outside option is removed. While parameter values shift under this alternative assumption, the key qualitative relationships remain consistent. For instance, the ratio of disutility from time spent driving versus using public transit is approximately 1.65 in our main model and 1.4 in this alternative specification. Both ratios confirm commuters' preference for passive transit over active driving, though the magnitude differs slightly between specifications.

Figure A.1 evaluates the fit of the model and is analogous to Figure 3; it compares the predicted share of public transit users according to the no-outside option model against corresponding shares from auxiliary data.[40] The predictions from our main model align considerably better with observed commuting patterns than those in Figure A.1, suggesting that accounting for alternative commuting arrangements improves model fit.

For completeness, we repeat the main counterfactual analyses under this alternative assumption. Table A.2 shows results analogous to Table 3, but with the outside transportation option removed. Generally, our findings are similar under this alternative assumption, although the impacts of SDVs are somewhat muted when the housing supply is inelastic. Moreover, Figure A.2—analogous to Figure 4—shows that driving directly displaces transit

[40] For more details on the construction of the figure, see Section 4.4.

ridership for most distance bands.

Finally, Table A.3 displays the impacts of counterfactual transit improvements both with and without SDV technology. These results align with our findings in the main text: transit improvements generate similar incremental revenue effects regardless of SDV availability, but such improvements cannot fully mitigate the increased vehicle miles traveled resulting from SDV adoption.

A.3 Congestion Zone Distances

We calculate arc distances to estimate the relative travel lengths between origin-destination pairs within each concentric circle. This approximation serves as a useful simplification for our congestion model that connects road usage patterns to average travel speeds across the circle band.[41] The procedure entails (1) constructing a great-circle path from the residence to either the workplace or commuter rail station (in cases of public transit usage), and (2) computing the length of path segments falling within each concentric zone.

For the innermost circle, the distance measure $M_{h,w,1}$ is defined by the length of the segment contained within the initial boundary, depicted by the red line in Figure A.3. For subsequent circles, we derive the distance by computing the differential between the current and preceding circles, expressed as:

$$D_{h,w,c} = M_{h,w,c} - M_{h,w,c-1} \tag{16}$$

where $D_{h,w,c}$ represents the distance within circle c exclusive of inner circles, $M_{h,w,c}$ denotes the total measured distance within circle c, and c indicates the circle index (with $c = 1$ corresponding to the innermost circle).

To implement this approximation of $M_{h,w,c}$ on a spherical surface, we first compute great-circle distances between (1) residence and workplace locations and (2) residence and the concentric circles' center, as well as the azimuths. We then determine the cross-track distance—defined as the minimum distance between the circles' center and the great-circle path

[41] We calculate distances between vehicle trip starting points and destinations using "as the crow flies" measurements. Rather than using purely straight lines, we employ arc distances to account for Earth's curvature, providing more accurate approximations of travel distances.

connecting residence and workplace—using the following spherical trigonometric relation:

$$d_x = \arcsin(\sin(d_{h0}) \times \sin(az_{h0} - az_{hw})), \tag{17}$$

where d_{h0} represents the great-circle distance from residence to circle center, az_{h0} denotes the azimuth from residence to circle center, and az_{hw} represents the azimuth from residence to workplace. The along-track distance from residence to circle boundary is subsequently derived through:

$$d_{at} = \arccos\left(\frac{\cos(d_{h0})}{\cos(d_x)}\right), \tag{18}$$

where d_{h0} denotes the distance from residence to circle center. The half-chord length of the intersecting line segment is computed as:

$$d_{half} = \arccos\left(\frac{\cos(d_r)}{\cos(d_x)}\right), \tag{19}$$

where d_r represents the circle's radius. This represents half the distance of the red line in Panel A of Figure A.3.

We compute the exit point of the residence to work path by projecting a distance of $d_{at} + 2 \times d_{half}$ along the trajectory. The total path distance within the circle $M_{h,w,c}$ is then derived through a systematic evaluation of geometric conditions, including: (1) whether residence, workplace, or both fall within circle c, and (2) whether any portion of the great circle path connecting home and workplace falls within the circle's boundary (otherwise the distance within the circle is zero). Figure A.3 illustrates various scenarios. Approximated driving miles exclusive to each circle are then calculated according to Equation (16).

A.4 Appendix Tables and Figures

Table A.1: Model Estimates – No Outside Option for Transportation

	Households Earning:	
	≤ 3,333 per month	> $3,333 per month
1(Involves Public Transit)	0.423	-0.4071
	(0.021)	(0.017)
Transfer Count	-0.060	-0.092
	(0.013)	(0.009)
Driving Duration	-0.058	-0.055
	(0.000)	(0.000)
Transit Duration (Non-Auto)	-0.042	-0.039
	(0.001)	(0.001)
Price ($)	-0.060	-0.070
	(0.001)	(0.001)

Notes: This table reports model estimates of the preference parameters for location and transportation characteristics from the demand model described in Section 4, when not allowing for an outside option. The model is estimated separately for low-income households (left column, with incomes at $3,333 or below) and high-income households (right column, with monthly incomes above $3,333). Standard errors in parentheses.

Table A.2: Simulated Outcomes of SDV Adoption – No Outside Option for Transportation

		Self-Driving Vehicles (SDVs)			
		Dynamic Congestion Effects			Congestion from Status Quo
	Status Quo	Persistent Location Choices	Fixed Housing Supply	Flexible Housing Supply	
Transportation Outcomes					
E[Miles Driven (one way)]	15.36	15.64	16.46	21.28	26.83
	(0.43)	(0.44)	(0.46)	(1.26)	(1.25)
E[Distance From Work]	16.28	16.28	17.13	21.96	27.37
	(0.43)	(0.43)	(0.46)	(1.26)	(1.24)
E[Public Transit Revenue Per Worker/Day]	1.02	0.88	0.89	0.90	0.83
	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
E[Public Transit Revenue Per Rider/Day]	4.68	4.39	4.43	4.42	4.15
	(0.02)	(0.01)	(0.01)	(0.02)	(0.03)
Percent Taking Public Transit	21.69	20.12	20.06	20.46	19.90
	(0.27)	(0.24)	(0.22)	(0.24)	(0.22)
$\frac{\textbf{Speed with Congestion}}{\textbf{Free-Flow Speed}}$					
Inner Ring	0.57	0.55	0.52	0.50	0.57
	(0.01)	(0.01)	(0.00)	(0.02)	(0.01)
Middle Ring	0.62	0.62	0.60	0.52	0.62
	(0.02)	(0.02)	(0.00)	(0.03)	(0.02)
Outer Ring	0.90	0.90	0.90	0.81	0.90
	(0.07)	(0.07)	(0.01)	(0.12)	(0.07)

Notes: This table shows estimated driving, location, and public transit usage outcomes in the baseline case without SDVs (Column 1) as well as in counterfactual scenarios when cars are autonomous, when we do not allow for an outside option. Columns 2 through 4 report outcomes when changes in miles driven affect driving speeds. These columns represent different assumptions about the housing supply: Column 2 maintains each individual's original probability of residing in any given census tract, Column 3 allows residential relocation but keeps the total housing supply per census tract fixed, and Column 4 assumes a perfectly elastic housing supply. Finally, Column 5 also assumes perfectly elastic housing supply but keeps traffic speeds at the status quo level. Standard errors in parentheses are calculated by parametric bootstrap, accounting for uncertainty in both consumer preference parameters and congestion model parameters.

Table A.3: Impacts of Public Transit Improvements – No Outside Option

	Miles Driven		Transit Revenue	
	No SDVs	SDVs	No SDVs	SDVs
Status Quo Transit	0.00%	38.55%	0.00%	-11.01%
	(0.00)	(5.03)	(0.00)	(0.63)
20% Onboard Travel Time Reduction	-1.99%	36.48%	21.90%	8.36%
	(0.07)	(4.97)	(0.17)	(0.88)
40% Onboard Travel Time Reduction	-4.37%	33.66%	50.10%	33.39%
	(0.10)	(4.92)	(0.44)	(1.23)
60% Onboard Travel Time Reduction	-7.51%	30.36%	86.81%	65.97%
	(0.15)	(4.84)	(0.87)	(1.72)
80% Onboard Travel Time Reduction	-11.26%	26.21%	134.96%	108.91%
	(0.23)	(4.73)	(1.50)	(2.44)
100% Onboard Travel Time Reduction	-16.03%	20.89%	198.45%	166.03%
	(0.29)	(4.66)	(2.47)	(3.53)

Note: This Table shows simulated percentage changes in miles driven (left two columns) and transit revenue (right two columns) for varying percent improvements in the speed of transit trains, relative to status quo transit speeds with no SDVs. Improvements are applied only to onboard (seated) travel time, while waiting times and the disutility associated with transfers remain fixed at baseline levels. Columns 1 and 3 report outcomes when there are no SDVs, and Columns 2 and 4 show outcomes, relative to the status quo, in the presence of SDVs. Standard errors in parentheses are calculated by parametric bootstrap, accounting for uncertainty in both consumer preference parameters and congestion model parameters.

Table A.4: Impacts of Public Transit Stops Improvements

	Miles Driven		Transit Revenue	
	No SDVs	SDVs	No SDVs	SDVs
Status Quo Transit	0.00%	40.90%	0.00%	-12.91%
	(0.00)	(1.81)	(0.00)	(0.68)
20% Transfer Hassle Reduction	-0.71%	40.26%	12.47%	-2.26%
	(0.07)	(1.80)	(0.47)	(0.82)
40% Transfer Hassle Reduction	-1.63%	39.47%	27.33%	10.39%
	(0.08)	(1.77)	(1.09)	(1.17)
60% Transfer Hassle Reduction	-2.42%	38.65%	45.12%	25.67%
	(0.11)	(1.75)	(1.88)	(1.75)
80% Transfer Hassle Reduction	-3.61%	37.55%	67.02%	44.33%
	(0.15)	(1.73)	(2.90)	(2.56)
100% Transfer Hassle Reduction	-4.77%	36.27%	94.33%	67.79%
	(0.18)	(1.68)	(4.25)	(3.64)

Notes: Simulated changes in miles driven (left two columns) and transit revenue (right two columns) for varying percent improvements in the hassle cost of changing trains, relative to status quo transit speeds with no SDVs. Columns 1 and 3 report outcomes when there are no SDVs, and Columns 2 and 4 show outcomes, relative to the status quo, in the presence of SDVs. Standard errors in parentheses are calculated by parametric bootstrap, accounting for uncertainty in both consumer preference parameters and congestion model parameters.

Figure A.1: Model Predictions & ACS Survey Data – No Outside Option for Transportation

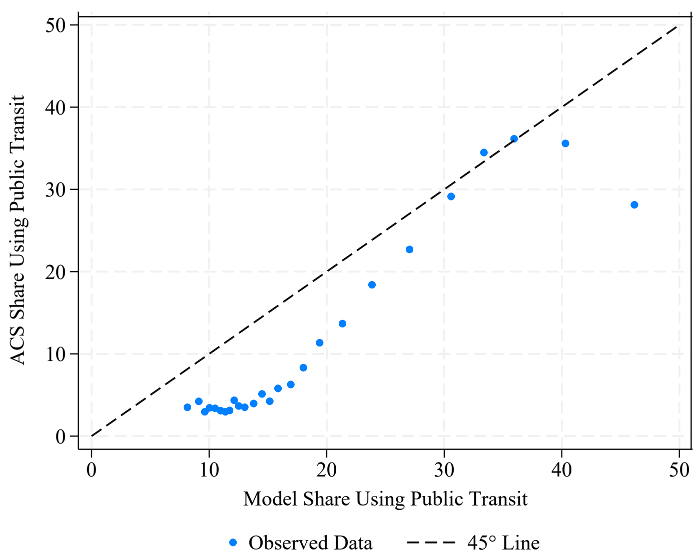

Notes: This figure compares the model's predicted census tract-level share of workers using public transit (x-axis) with 5-year estimates based on survey data from the 2014-2018 American Community Survey (ACS) (U.S. Census Bureau, 2018), when not allowing for an outside option. Each point represents 50 combined census tracts, grouped based on the predicted percentage of public transit usage. The 45-degree red dashed line represents perfect matches between the model's prediction and the American Community Survey's estimate.

Figure A.2: Transportation Mode Choice Changes – No Outside Option for Transportation

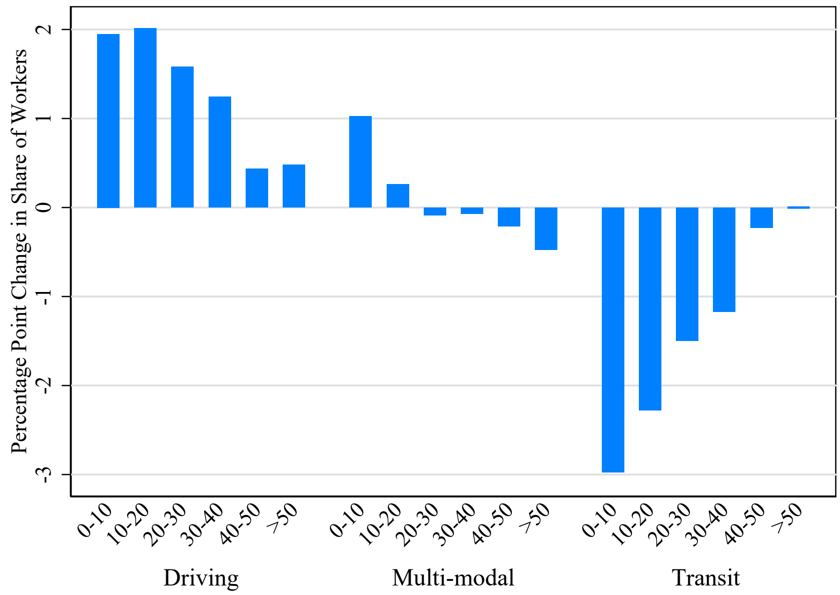

Notes: The figure shows changes in the shares of workers using each of three modes of transportation when SDVs are introduced and home locations are persistent, for different distance ranges, from 0 to 10 miles, to over 50 miles.

Figure A.3: Illustration of Great-Circle Paths Within Concentric Zones

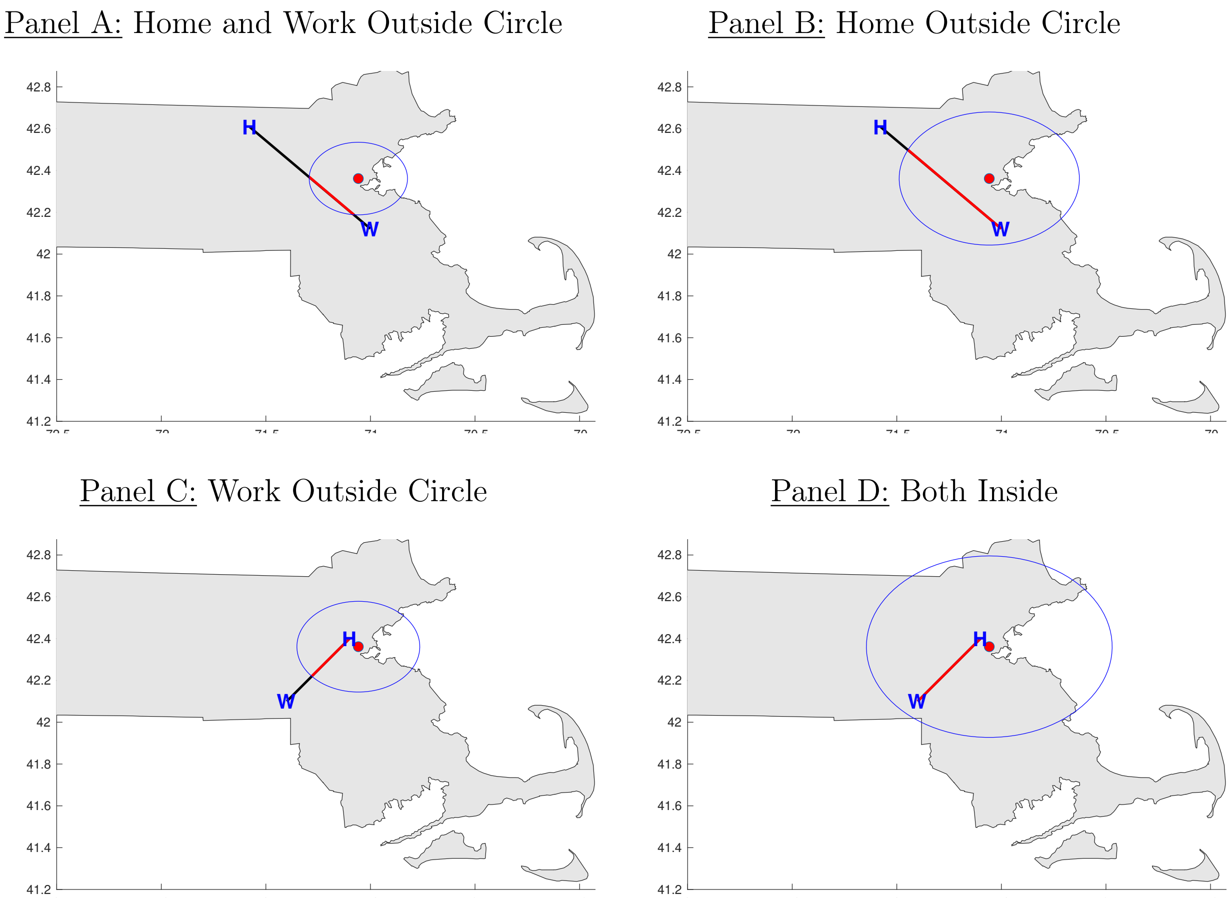

Note: The diagram centers on Boston City Hall, marked by a red circle. This serves as the focal point for our concentric circle model. A hypothetical worker's residence and workplace are indicated by 'H' and 'W' respectively. The worker's straight-line commute path is depicted by a black line linking these two points. The segment of this route that falls within the circle is highlighted in red, illustrating the portion of the journey affected by the modeled traffic congestion.