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

Imke Reimers
Cornell University

Benjamin Reed Shiller
Brandeis University

March 14, 2025

Abstract

This paper examines the impact of self-driving vehicles (SDVs) on commuter decisions and the efficacy of public transportation investments. We estimate consumer preferences for home locations and commuting options using detailed data from the Greater Boston area, focusing on the trade-offs between driving and public transit. Our results suggest that SDVs, by reducing the disutility of driving, could increase miles driven by 40 percent and decrease public transit use and revenue by about 10 percent. While improvements in public transit – such as faster trains and reduced transfer costs – still meaningfully raise transit revenues, the effects on mileage are modest compared to the impacts of SDVs. These findings highlight the challenges and potential of public transportation investments in the era of autonomous driving technologies.

1 Introduction

Investment in long-lived transportation infrastructure requires planning, and the development of self-driving vehicles (SDVs) – such as Mercedes' recently released hands-off technology – may change spatial patterns of driving and transit use by decreasing the opportunity cost of driving. In this paper, we examine the interplay of these demand-side innovations and improved public transportation options in commuter decisions. Specifically, we ask (1) how autonomous driving technology will affect residential choice and commuting mode decisions, and (2) how these changes affect the returns to investments in public transit.

To evaluate the effects that self-driving vehicles (SDVs) will have on commuter decisions and on the efficacy of investment in public transportation, we first need to understand 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 choice, we flexibly jointly estimate demand for housing location and transportation choices. Specifically, we model transportation mode demand as a function of its inherent attributes: active driving time, passive travel time (e.g., on a train), waiting time between trip segments, number of stops, and cost. Although we do not observe transportation mode choice directly, we infer preferences for these characteristics through the individuals' residential location decisions. The degree to which each transportation mode feature influences home location choice indicates the strength of preference for that feature. The variation of work locations across our sample allows us to disentangle preferences for transportation characteristics from intrinsic preferences for specific residential areas.

Our estimates for both low-income and high-income individuals suggest that characteristics of the commuting options between home and work locations (e.g., driving time, transit time, count of layovers) significantly influence their housing choices. Consumers' choice of

residential locations that are relatively easy to reach by public transit identifies their relative disutility of driving time over time spent passively on public transit. We interpret these estimates as evidence that consumers value having their attention free for other activities during commuting time. Comparisons of our estimated commuter rail ridership with aggregate and location-specific train ridership information from separate data sources validate our results.

Armed with these preference estimates, we evaluate both main research questions. First, we examine how advances in driving technology such as SDVs will affect consumer commuting decisions, including their miles driven, as well as 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. Holding fixed the public transit network, We find that the widespread adoption of SDVs significantly raises miles driven, by 38 percent. This includes 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 at the same distance. In addition to raising driving miles, the replacement of all conventional cars with SDVs also reduces public transit revenue, by about 8 percent; and the share of people using public transit decreases from 12.4 percent to 11.4 percent. Partial transition from conventional vehicles to SDVs produces effects that are proportional to the rate of adoption.

While the arrival of SDVs decreases public transit use overall, it is still unclear whether it changes the marginal returns to investment in public transit improvements. We examine these impacts by simulating transit improvements, both in scenarios with and without SDVs. Transit improvements lower car miles and raise day-to-day transit revenue by similar amounts with and without SDVs, suggesting that investments in transit continue to yield desired returns.

Finally, we ask whether the same transit improvements can offset the effects of SDVs on driving and transit revenue. For car mileage, the impacts of transit improvements are small compared to the direct effect of SDVs. For example, reducing on-transit travel times by 40 percent would reduce car mileage by only 2 percent – much less than the 40 percent increase in miles driven due to SDVs. By contrast, public transit improvements can generate sufficient revenue to offset the losses caused by SDVs. Even a 20 percent speed improvement raises transit revenue enough to offset the decrease in revenue arising from SDVs. Our results indicate that improvements in public transit can still raise revenues and yield material benefits for riders.

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 presents evidence, based on auxiliary datasets reporting commuter rail ridership, of our model's fit. In Section 5, we use the estimates from the model to run counterfactual simulations with SDVs as well as improvements in public transportation. Finally, Section 6 concludes with implications for current investment decisions and long-term policies.

2 Background

The paper compares the effects of two separate changes in commuting options to work, the widespread introduction of a new driving technology in SDVs, and improvements to public transit options. These counterfactual changes are 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

People living in one location and working in another are faced with various options for commuting. In the United States, two of the most commonly used options including driving and – where possible – taking public transportation. Driving continues to be the top choice with about 72 percent of commuters reporting using their car to get to work.[1] Public transportation, at about 13 percent, is the second choice, just ahead of walking (12 percent) and riding a bike (10 percent). These numbers may somewhat overstate the role of cars and understate the role of public transit in areas where public transportation is a more viable option. For example, in metropolitan areas in the Northeast, the share of commuters using public transit exceeded one third in 2019.[2] 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

While public transit is less developed and less used in the U.S. relative to other countries, public transportation takes on a sizable – and growing – share of total government spending across the US.[3] The congressional budget office reports annual investments for public transit

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

[2] See https://www.census.gov/content/dam/Census/library/publications/2021/acs/acs-48.pdf.

[3] See https://www.statista.com/outlook/mmo/shared-mobility/public-transportation/worldwide#global-comparison.

projects from 1956 to 2014.[4] Over this time period, the share of total 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.[5] 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, for a 29-fold increase in real dollars.

These trends remain consistent in current years: In April 2024, the US Department of Transportation announced it would make $20.5 billion available to support public transit investment projects throughout the country as part of "history's biggest-ever federal investment in transit – trains, buses, ferries, and more – so people can reliably and affordably get to wherever they need to be," according to U.S. Transportation Secretary Pete Buttigieg.[6]

The city of interest in our study, Boston, is serviced by the Massachusetts Bay Transportation Authority (MBTA), which 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 MBTA's mission, affirmed by its 2023 Strategic Plan, is to "serve the public by providing safe, reliable, and accessible transportation."[7] Each day, the MBTA commuter rails alone transport an average of about 94,000 riders, and the MBTA recently purchased additional commuter rail cars to match rising demand.[8] Average weekday use across all public transportation in the Boston area is reported at 749,199 people

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

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

[6]See https://www.transportation.gov/briefing-room/biden-harris-administration-provide-205-billion-fed and see https://www.transit.dot.gov/funding/grants/fy24-fta-bus-and-low-and-no-emission-grant-awards for a list of funded projects.

[7]See https://www.mbta.com/news/2024-06-11/mbta-board-approves-capital-investment-plan-fiscal-years-20

[8]See https://www.mbta.com/performance-metrics/ridership-the-t and https://www.mbta.com/news/2024-08-14/mbta-purchases-additional-new-commuter-rail-cars-improve-service-ridership-grows.

in August 2024.

The costs associated with running the local public transit system are far from negligible. The systems requires estimated expenditures of about $25 billion to address aging tracks, stations, and other infrastructure issues.[9] Any improvements 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.[10] Together, these costs would amount to about $8,000 per Massachusetts resident, including many of whom do not use (or live near) public transit. These high costs raise questions about whether these investments will retain their intended benefits in the presence of SDVs.

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.[11] Many 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 people travel between places.

The Society of Automotive Engineers (SAE) has established a widely recognized frame-

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

[10] 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.

[11] 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.

work 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 "Full Self-Driving" provide a prominent example of Level 2 cars. These cars allow the system to manage both steering and acceleration/deceleration under certain conditions but require constant driver supervision. One step beyond this, 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). 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. Another step further, Waymo, a subsidiary of Alphabet, offers Level 4 autonomous rides without a driver in geographically bounded areas in Los Angeles CA, San Francisco CA, and Phoenix AZ, and it is soon expanding to Austin TX, Atlanta GA, and Miami FL.[12] 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 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 will likely reach mid-range vehicles within 5 years and economy models within a decade.[13]

With increasingly affordable autonomous driving technology, SDVs may soon be widely

[12]See https://waymo.com/about/, accessed January 14, 2025.

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

available to the public. However, the rate of adoption will depend on two key factors: how much consumers value using commute time for other activities versus 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."

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 (Cramer and Krueger, 2016; Chen et al., 2019; Frechette et al., 2019; Guda and Subramanian, 2019).[14] However, these benefits are accompanied by increased congestion and environmental costs (Tarduno, 2021).

While ride-hailing services are well studied, the scarcity of empirical studies on Self-Driving Vehicles (SDVs) specifically reflects their limited deployment. 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 poten-

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

tially enable higher speeds with reduced vehicle spacing, increasing road capacity (Karbasi et al., 2023; Lu et al., 2020). This gain in efficiency may not translate to faster travel times if SDVs raise overall driving; and survey evidence suggests that the perceived cost of travel time could decrease by 18–32 percent (Zhong et al., 2020). Hence, SDVs may 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 transport 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 12-25% 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 benefit-cost calculations.

Finally, our estimation of home location choices builds on literature examining how transportation innovations affect urban form and travel patterns.[15] 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; Barwick et al., 2024). 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 many analyses focus on market-specific equilibrium effects, some (Almagro et al., 2024; Durrmeyer

[15] See Weisbrod and Reno (2009) for an overview.

and Martinez, 2024; Glaeser and Kahn, 2010) also emphasize the importance of considering environmental externalities. Our contribution to this literature is to explore how technological advances will reshape the effectiveness of infrastructure investments.

Among the existing literature, Barwick et al. (2024) represents the closest methodological parallel to our work. Like Barwick et al. (2024), we simultaneously estimate transit preferences and housing choices. However, our approaches diverge in two key aspects. First, while they derive transit preferences directly from observed choices, we infer transportation preferences through residential location decisions. Second, their housing demand estimates incorporate local characteristics explicitly, whereas our approach uses location-specific mean intrinsic values, abstracting from location characteristics. Although our approach sacrifices some precision by not utilizing direct transportation choice data, it fits home locations choices perfectly and offers meaningful advantages in terms of data requirements, allowing broad application in settings where systematic transit mode choice data are not available.[16] We demonstrate in Section 4.4 that this simplified method maintains good fit.

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 driving, public transportation, and a third option of driving to public transit access points. We obtained driving times through the OpenRouteService API, and we collect public transit travel times and routes from the Massachusetts Bay Transportation Authority (MBTA) Trip

[16]The home and work locations data used in this paper are available for all geographic areas in the US.

Planner. We supplement these data with commuting and parking prices.[17]

The LEHD Origin-Destination Employment Statistics (LODES) dataset provides information on home and work locations by census tract pair.[18] 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.[19] To make the 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) and the 1301 home census tracts in a radius of 75 miles around Boston's city center.[20] Our sampling retains about 35% of all workers in the LODES Massachusetts dataset: Restricting work locations to fall within 20 miles of Boston retains 56.6% of workers in the dataset, and restricting to the 100 most worked locations retains 61.9% of the remaining workers.

To systematically analyze commuting patterns and preferences for different transportation characteristics, we collected travel times between census tract centers for three transportation modes: driving, public transit, and a hybrid option of driving to public transit.[21] We obtained driving times and miles from the OpenRouteService API, which accounts for typical congestion levels; and we calculate public transit times using MBTA's Trip Planner data.[22] For public transit and hybrid routes, we specified arrival times of 9:00 AM on Wednesdays to reflect typical commuting conditions. We decomposed each route's total travel time into active travel time and passive time (waiting at transfer points), while also recording the number of transfers required. We calculated travel times using the geographic

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

[18] See https://lehd.ces.census.gov/data/lodes/. We use version **7**, which uses 20**10** census block definitions and contains data from 2018.

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

[20] We also allow consumers to choose to live in locations beyond 75 miles from Boston. In the model, these locations constitute the outside good.

[21] For the hybrid option, we computed times using the closest stop on each commuter rail line and selected the fastest resulting route.

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

centers of census tracts as origin and destination points.

We supplement the commuting time information with data on round trip transportation prices. Specifically, we use twice 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 (round trip), which is the IRS-suggested variable cost for 2018, plus the destination parking cost inferred from the closest public parking lot.[23] Public transit prices seem to be set exogenously without consideration of cross-location variation in demand. Figure 2 plots transit prices for different commuting zones, in 1989 vs. 2024. The prices lie almost exactly on a straight line, suggesting no relative price changes over a 35-year period.

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 for using only public transportation.[24] 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). Driving, on the other hand, tends to be much faster: On average, people can reach their work location by car in 28.9 minutes (median = 23.6). Only 0.2 percent of census tract pairs have faster public transit connections than driving connections. Despite this rarity of transit being the quicker option, many people still choose public transit. 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 segments across all census pairs – not weighting by chosen locations – are significantly longer than those reported in Table 1, at about 101 minutes (vs.

[23]The medical care transportation cost per mile is considered the variable cost of driving. See https://crsreports.congress.gov/product/pdf/IN/IN12320.

[24]The number of observations is lower because not all census tract pairs can be connected using only public transit.

74) and 2.7 segments (vs. 2.2) for public transit connections, and 50 minutes (vs. 29) for driving connections. The fact that average commuting times that account for the frequencies of observed home/work pairs are much shorter than the raw averages suggests that people value shorter and more convenient commutes.

Figure 1 provides additional suggestive evidence to support this takeaway. 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.[25] 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 the residential choice model – given work locations – described here.[26] The basic premise is that, ceteris paribus, individuals exhibit a higher propensity to reside in areas where appealing transportation options to their workplace are convenient. This allows us to infer people's transportation preferences by examining where they choose to live without needing information on their daily travel decisions. In the next section, we will use these estimated preferences to examine both location and transportation choices when transportation mode characteristics.

[25] Given that the rail and highway networks were established before much of the population increase suggests that homes were built near highways, not the other way around.

[26] 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.

4.1 Model

The conditional indirect utility that consumer i in income group g (from set $G = \{low, high\}$) receives from transportation mode c (from set $C = \{car, transit, both, 0\}$, where 0 denotes the outside option of biking, walking, or working from home) 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, and the outside option 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 choices is the inclusive value λ:

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

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

The conditional indirect utility of residential location choice h is 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}):

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

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

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

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:

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

where $\mathbb{s}_{gh}$ 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 corresponding δ_{gh}, the log likelihood of workers' home location choices equals:

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

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

Here, we explain how people's 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 spacial variation in employment loca-

tions. Whereas location amenities are common across all consumers in group g and location h, transportation attributes – transportation mode-specific travel time and costs, and the number of transfers – vary based on an individual's workplace. 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.

A few specific examples can clarify the intuition. A consumer working at Brandeis University (in Waltham) can use the Brandeis stop on the Fitchburg commuter rail line, making commuting by public transit fast and convenient for any home locations on the same line. Thus, a Brandeis employee living in Concord could arrive at Brandeis in about 18 minutes, whereas driving during rush hour would typically take between 16 and 24 minutes.[27] By contrast, public transit options from locations *not* on the same line would require both longer travel times and more transfers. A consumer working at Brandeis and living in Dedham, for instance, would have the same driving commute as the Concord resident, but their public transit commute would take 72 minutes and involve two transfers. Holding amenities constant, if people working at Brandeis are more likely to live in Concord than in Dedham, our model would interpret this as consumers valuing convenient and fast public transportation and disliking transfers. If, on the other hand, people at Brandeis were about equally likely to live in the two locations, the model would determine driving time as the more important determinant and therefore as the preferred commuting option with lower commuting disutility.

Variation in trains 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

[27]Travel times taken from Google Maps for arriving on a Friday at 9am.

per hour.[28] By contrast, the Green Line, despite sharing two stops with the Orange Line, moves much more slowly, at an average of only 10 miles per hour and some segments crawling as slowly as 3 miles per hour.[29] Higher propensities to live near the orange line, at similar distances from workplaces, indicates commuters value the 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 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 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 individual effects of characteristics, we sidestep the price endogeneity problems that typically arise when decomposing these effects..[30]

Turning to transportation costs, we argue that mode-specific prices are plausibly exogenous to local residential choices. This exogeneity is particularly clear for private vehicle usage, where the primary cost components are determined outside the local housing mar-

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

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

[30] Remaining agnostic about the roles of specific location characteristics allows us to estimate location preferences perfectly, but it raises some concerns about potential equilibrium changes in location characteristics when people move in counterfactual scenarios (Baum-Snow and Kahn, 2000; Tsivanidis, 2023; Barwick et al., 2024). While these concerns are somewhat alleviated by the fact that most migration is to locations farther from the city center (that are less space constrained), we also provide bounds for the overall effects by introducing counterfactual analyses in which consumer relocation is restricted.

ket. Gasoline prices, while varying with commute distance through fuel consumption, are set in national markets and thus insulated from Boston-specific demand patterns. Additionally, workplace parking costs depend on the work location, and hence are invariant to an individual's choice of residential location.

Public transit fares in our context also plausibly satisfy exogeneity conditions. 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. Figure 2 illustrates this point by comparing the zonal pricing structure of the commuter rail in 2024 to the prices in 1989.

The above evidence suggests that our modeling choices and the institutional details address obvious endogeneity concerns.

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). 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, a minute of driving creates much more disutility than a minute on public transit in both income groups. High-income consumers, for example, equate one driving minute to 1.591 minutes on public transit (-0.070/-0.044), indicating willingness to spend 59.1% longer

on transit vs. driving. We interpret this lower public transit time disutility as stemming from passengers' ability to multitask, similar to what autonomous vehicles allow. We leverage this in counterfactuals by assuming time spent in self-driving vehicles (SDVs) generates similar utility to public transit time.

Second, the mode-specific indicator coefficient reveals a significant fixed disutility associated with public transit utilization, with pronounced heterogeneity across income groups. Converting these estimates to driving-time equivalents, we find that the fixed disutility of transit usage corresponds to 1.4 minutes of driving time for lower-income households (-0.111/-0.078) but rises to 6.7 minutes (-0.470/-0.070) for higher-income households.

Similarly, price sensitivity also exhibits expected heterogeneity across income groups. While most coefficients are similar in magnitude between groups, the price coefficient for low-income households (-0.018) is twice that of high-income households (-0.009). 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 4.1 minutes of transit time for low-income users (0.210/0.051) and 4.8 minutes for high-income users (0.210/0.044). Since the expected transfer time is already included in total trip duration, this disutility reflects other factors like schedule uncertainty or required attention during transfers. These findings carry implications for autonomous vehicles: by offering point-to-point service, SDVs eliminate transfer penalties entirely while allowing multitasking, thus creating an inherent advantage over multi-segment transit routes.

4.4 Empirical Validation

The model infers preferences for transportation mode characteristics without incorporating actual transportation choices. Consequently, the model predictions may not reflect actual transportation mode choices. Here, we use external survey data to assess the validity of the implied transportation preference parameters and the resulting choices, in three analyses.[31] 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 estimates of transit ridership patterns across census tracts. Third, we compare the model's predicted use of multi-modal commuting options with survey data.

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 12.2% public transit usage rate, while the ACS data indicates an 11.1% rate.

Next, we compare model-predicted public transit usage with usage reported in the 2014-2018 5-year ACS data. Specifically, we create groups of 50 tracts each, grouped according to our model's predicted public transit usage share; and for each of these groups, we compute the worker population-weighted average percentage choosing public transit. Figure 3 shows the share using public transit in each group according to ACS survey data (y-axis) against the corresponding model predictions (x-axis). The analysis shows that the model predictions track external data predictions relatively closely, although there are some discrepancies in census tracts with high predicted public transit usage.

[31] 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.

Finally, we test our model with external data on station access behavior among commuter rail passengers. Our model predicts that 61.4% of commuter rail commuters use private vehicles to access stations. This prediction aligns well with empirical evidence from a 2008-2009 MBTA passenger survey (Humphrey et al., 2010), which found that 65.3% of passengers reached stations via private motorized modes (including personal vehicles, carpools, and private shuttles).

This close alignment of the model predictions and external data on these three dimensions suggests that the model reflects transportation choices despite not using this information directly in the estimation. This suggests that residential and workplace location data alone can be used to analyze transit usage patterns and counterfactual policies, without requiring transit ridership survey data.

5 Equilibrium Effects of SDVs and Investment

Our objective is to compare the transportation outcomes resulting from self-driving vehicles (SDVs) with those from investments in traditional public transit. We achieve this through counterfactual simulations, modeling SDVs as having the same road network capabilities as traditional vehicles and the same travel time disutility as public transit. Each simulation accounts for congestion changes and establishes a new joint equilibrium of congestion, home location, and transportation mode choices.

To separate the effects of transportation mode choice and home location choice, we examine different 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. Additionally, two other counterfactuals are used to provide intuitive bounds for the effects when relocation is permitted. In one case, home relocation is allowed,

but the local housing stock remains fixed at current levels. In the other case, 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. 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 equilibrium by iteratively updating zone-specific speeds until they produce residential and commuting patterns that generate VMT levels consistent with those speed assumptions. This subsection describes the congestion model, parameter calibration, and equilibrium solution method.

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. Using current VMT and travel speeds from each zone, we calibrate zone-specific free-flow speeds (v_{fz}) and road capacities (k_{0z}), which allows us to map zone-level VMT (k_z) to the resulting average traffic speeds (v_z) in counterfactual simulations. We find v_{fz} and k_{0z} in three steps: We first estimate the travel speeds v_z and v_{fz} in the current environment, then we infer status quo VMT from our model, and finally we use these parameters to calculate the corresponding road capacities k_{0z} via Equation (7).

To estimate current traffic speeds we collect travel times for 100 home-work location pairs from Google Maps on Tuesday at two different times of day: first, to arrive at 9am (to capture v), and second, when leaving at 3am (to capture v_f, the free-flowing traffic speed). 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 $\left(\frac{1}{v_z}\right)$ as coefficients from a regression of t_ℓ on the zone-specific distances $d_{\ell z}$. To obtain estimates of the zone-specific traffic speeds without traffic (v_{fz}), we estimate this with off-peak traffic data (3am); and to estimate speeds with traffic, we use data from 9am.[32]

[32]To address concerns that travel time may depend on route, which may change in counterfactual environments, we measure both distance and velocity as the components along the straight-line distance. See Appendix A.3 for details.

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

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, we use an iterative process. We start with an initial guess of traffic speeds, $\tilde{v}_z$, for each traffic zone. Based on these speeds, we simulate home choices and transportation choices, which yield a corresponding VMT in each zone (k_z). Using these distances, 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.

The estimated speeds achieved during rush hour relative to free-flowing traffic (v/v_f) are 57 percent in the inner ring (closest to Boston center), 65 percent in the middle ring, and 97 percent in the outer ring. 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 in the long run when SDVs are introduced and fully adopted. 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, with varying assumptions regarding housing relocation. Columns

2 through 4 endogenize traffic speeds to respond to congestion levels. Specifically, the second column isolates the effects of SDVs on transportation mode choice, excluding any impacts from residential relocation;[33] and the third and fourth columns introduce scenarios where relocation is permitted. In the third column, the income-type-specific housing supply at the census tract level is fixed at status quo levels, while the fourth column allows for new construction at prices consistent with existing housing.[34,35] The 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).

First, consider the impacts of SDVs on miles driven and distance from work – 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 24% (9.53 to 11.80 miles) when home locations are fixed, to 37.8% (9.53 to 13.13 miles) when the housing supply is perfectly flexible. The difference in vehicle usage effects arises from the result, in the second row, that consumers choose to move further from their work locations when driving is made less burdensome. We illustrate these moves in Figures 4 and 5. In Figure 4, we model the changes in housing prices that would keep the

[33] 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.

[34] The simulation in Column 3 adjusts home location probabilities while restricting housing supply to preserve baseline tract-level resident counts. This is accomplished by adjusting type- and home-specific inclusive values ($\tilde{\delta}_{gh}$). The procedure resembles that in footnote 33, except that the inclusive values are now fixed at the group- and home-specific status quo level.

[35] For computational tractability, we model separate housing markets for low- and high-income households. This segmentation reflects real-world policies, as exemplified by Massachusetts legislation that mandates municipalities to maintain a share of affordable housing units restricted to residents earning below 80% of the county's median income.

housing supply unchanged (see Column 3 of Table 3), and Figure 5 shows where people would move if housing prices remained constant (Column 4). Both figures reflect large increases in housing appeal in the outer suburbs of Boston.

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 13.7% (from \$0.80 to \$0.69 per day). This decrease arises from both a decrease in the revenue per rider (from \$6.53 to \$6.14) and a decrease in the share of workers who use public transit (from 12.2% to 11.2%). 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 people driving to transit access points – public transit revenue decreases by only 8.8%. The smaller revenue reduction occurs because the ridership decline is less severe and the revenue per rider almost completely recovers.

Figure 6 illustrates the changes in transportation mode choices based on the Column 3 specification, for workers living various distance ranges from their work locations.[36] 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.

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, but traffic speeds are held constant at baseline levels. This scenario suggests an even greater shift toward driving

[36]The figure looks almost identical when instead assuming perfectly elastic housing supply.

(raising miles driven by 66.0% relative to the status quo, and by 20.5% relative to Column 4), as well as home choices even farther from work locations, to an average distance of 18.44 miles. These changes also imply a sharper decline in public transportation use, with the share of users dropping to 11.14%.

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 $32.14 per workday when accounting for congestion, and up to $54.85 when keeping traffic speeds unchanged. The gains are largest for high-income consumers, whose CS rises by $43.35 (or $72.99 without congestion), whereas low-income CS rises by just $8.67 ($16.90). 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, though it may also introduce new sustainability concerns.

5.3 Rising Shares of SDVs

The above analysis assumes that SDVs fully replace non-autonomous cars. This is an extreme assumption, as some consumers may remain resistant to using fully autonomous vehicles (see Smith and Anderson, 2017). Here, we examine the effects from above when SDVs account for certain shares of the full fleet of cars, ranging from zero percent to 100 percent.[37] The left panel of Figure 7 shows how average miles driven change as we raise the share of SDVs, and the right panel shows changes in the transit revenue. The relationships between the share of SDVs and both outcomes are almost linear, with values bounded between those in columns 1 and 2 of Table 3. The same patterns also hold for the other outcomes reported

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

above. Moving forward, we continue to assume full penetration of SDVs, noting that the results are roughly proportional for lower take-up rates.

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 gains to usage and revenue? To address this question, we examine the effects of public transit improvements on revenue and VMTs, both in a world with and without SDVs. In the main analysis, we model improvements in the speed of individual trains (which we model as a reduction in the passive travel time in public transit). 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 during transfers). All results remain, although traffic stop improvements have slightly smaller effects.[38]

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 trains were to decrease by 40%, some consumers would switch to public transit and miles driven would decrease by 2.8% (se=0.??).

We find similar marginal impacts in a scenario with SDVs (Column 2). For example, improving transit speeds by 40% reduces VMT by about 2% ($\frac{137.85-134.98}{137.85} \times 100 \approx 2.1\%$). However, these decreases in miles driven pale in comparison with the impact of SDV adoption.

[38]These relative results – implying wait times are less salient or valued – are consistent with, for example, Buchholz et al. (2020), who show that consumer waiting time elasticities are only about a fourth the size of their price elasticities in ride hailing.

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 24.27% (se=2.??) higher than under the status quo without transit improvements or SDVs. Likewise, if SDV adoption were coupled with an elimination of transfers (with time on trains remaining unchanged) miles driven would be be 33.13% (1.??) above status quo levels. These findings reflect the persistence of mode choice, combined with large increases in mileage from moving further from one's place of work.[39]

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. Note from Figure 7 that the effects of SDVs are relatively linear in their adoption rate. This allows for back of the envelope calculations for a threshold penetration rate ϕ of SDVs 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% reduction in time between entry and exit stations would still be positive – more mileage – if SDV penetration exceeded 7.8%.[40] Even with 100% speed improvements, VMT would still exceed baseline levels for any penetration rate above 35.8%.

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% improvement in train speed would raise transit revenue by 28.9% without SDVs; and even with SDVs, revenues are 17.1% above baseline levels. Similar patterns ensue from reducing the hassle cost of transfers.

[39]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.

[40]With a 40% reduction in passive transit travel time, the overall change in mileage with SDV penetration rate ϕ equals $(1-\phi) \times -2.83 + \phi \times 34.98$. This expression exceeds zero when $\phi \geq 0.0748$.

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 40% raises revenues by 27.4%, only slightly below the effect of the same improvement in the absence of SDVs (28.9%).[41] 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 depends on the policymakers' 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 improvements. 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 something else. While transit improvements reduce vehicle miles traveled, these reductions are modest relative to the substantial increases induced by SDVs.

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, for example, Almagro et al., 2024; Barwick et al., 2024; Durrmeyer and Martinez, 2024; Kreindler, 2024). The consistency of investment effects across scenarios

[41] $\frac{17.09\%+8.12\%}{100\%-8.12\%} \approx 27.4\%$.

with and without SDVs suggests that such policy interventions could extend to a future with SDVs.

Ultimately, a mix of policies may be necessary to effectively shape urban 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 available of autonomous driving technologies suggest that it is 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 two main questions: will SDVs substantially change the way that we commute, potentially raising environmental concerns further; and if so, will investments in public transit lose their efficacy?

This paper provides a first answer to this question. Using data on home and work locations for a census of people working in the Boston area, along with information on travel time by various modes of transportation, we estimate a structural model of commuting and home location choice that allows us to model counterfactual scenarios in which SDVs are widely used, as well as improvements in public transportation. We find, first, that the widespread adoption of SDVs significantly changes both home locations and transportation choices: people will likely move further away from their work and become more likely to use cars for commuting, leading to increased congestion and lower public transit revenue.

Second, we find that while improvements in the public transit experience cannot reverse the increase in miles driven, they *can* raise public transit revenue if the improvements are large enough. These results suggest that while environmental concerns need to remain at the forefront of transit policy, public transit can remain an important and valuable means of transportation looking forward.

Our paper just captures 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 part of 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.

Baum-Snow, N. and M. E. Kahn (2000): "The effects of new public projects to expand urban rail transit," Journal of Public Economics, 77, 241–263.

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. Matějka, and T. Salz (2020): "The value of time: Evidence from auctioned cab rides," .

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.

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

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

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.

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.

Karbasi, A. H., B. B. Mehrabani, M. Cools, L. Sgambi, and M. Saffarzadeh (2023): "Comparison of speed-density models in the age of connected and automated vehicles," Transportation research record, 2677, 849–865.

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.

Lu, Q., T. Tettamanti, D. Hörcher, and I. Varga (2020): "The impact of autonomous vehicles on urban traffic network capacity: an experimental analysis by microscopic traffic simulation," Transportation Letters, 12, 540–549.

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.

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.

Zhong, H., W. Li, M. W. Burris, A. Talebpour, and K. C. Sinha (2020): "Will autonomous vehicles change auto commuters' value of travel time?" Transportation Research Part D: Transport and Environment, 83, 102303.

7 Tables and Figures

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

	Summary Statistics—Weighted by Home/Work Pair 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 ($)	52,132	8.74	9.14	3.40	0.00	60.40
Driving to Transit						
Duration (transit)	97,735	50.35	26.96	46.00	0.00	168.00
Duration (driving)	97,735	35.20	17.90	31.69	0.00	238.16
# transit segments	97,735	1.69	0.93	2.00	0.00	6.00
Price ($)	97,735	30.46	11.09	30.05	1.67	79.91
Driving						
Duration	105,770	28.86	21.03	23.61	0.00	236.14
Price ($)	105,770	25.20	15.32	25.84	0.00	83.35
Fastest Commute Mode						
Duration	105,770	28.86	21.03	23.60	0.00	236.14

Notes: 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 people observed in each home/work location pair. "Public Transit Only" describes all pairs for which commuting by only public transit is possible; "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: Model Estimates

	Households Earning:	
	$\leq$ 3,333 per month	$>$ \$3,333 per month
Constant	5.084	5.385
	(0.023)	(0.024)
1(Involves Public Transit)	-0.111	-0.470
	(0.028)	(0.019)
Transfer Count	-0.210	-0.210
	(0.031)	(0.020)
Driving Duration	-0.078	-0.070
	(0.000)	(0.000)
Transit Duration (Non-Auto)	-0.051	-0.044
	(0.001)	(0.001)
Price (\$)	-0.018	-0.009
	(0.001)	(0.001)

Notes: 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 in parentheses.

Table 3: Outcomes in the Status Quo vs. with Self-Driving Vehicles (SDVs)

		Self-Driving Vehicles (SDVs)			
		Dynamic Congestion Effects			Congestion
	Status Quo	Persistent Location Choices	Fixed Housing Supply	Flexible Housing Supply	Same as Status Quo
Transportation Outcomes					
E[Miles Driven (one way)]	9.53	11.80	12.47	13.13	15.82
	(0.13)	(0.13)	(0.11)	(0.23)	(0.27)
E[Distance From Work]	14.38	14.38	15.12	16.51	18.44
	(0.09)	(0.09)	(0.11)	(0.19)	(0.22)
E[Public Transit Revenue Per Worker/Day]	0.80	0.69	0.70	0.73	0.71
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
E[Public Transit Revenue Per Rider/Day]	6.53	6.14	6.30	6.45	6.33
	(0.04)	(0.03)	(0.03)	(0.04)	(0.03)
Pct Take Public Transit	12.24	11.24	11.16	11.39	11.14
	(0.29)	(0.25)	(0.24)	(0.26)	(0.24)
$\frac{\text{Speed with Congestion}}{\text{Free-Flow Speed}}$					
Inner Ring	0.57	0.52	0.50	0.48	0.57
	(0.01)	(0.02)	(0.00)	(0.02)	(0.01)
Middle Ring	0.65	0.58	0.57	0.52	0.65
	(0.02)	(0.02)	(0.00)	(0.02)	(0.02)
Outer Ring	0.97	0.93	0.92	0.91	0.97
	(0.05)	(0.07)	(0.00)	(0.08)	(0.05)

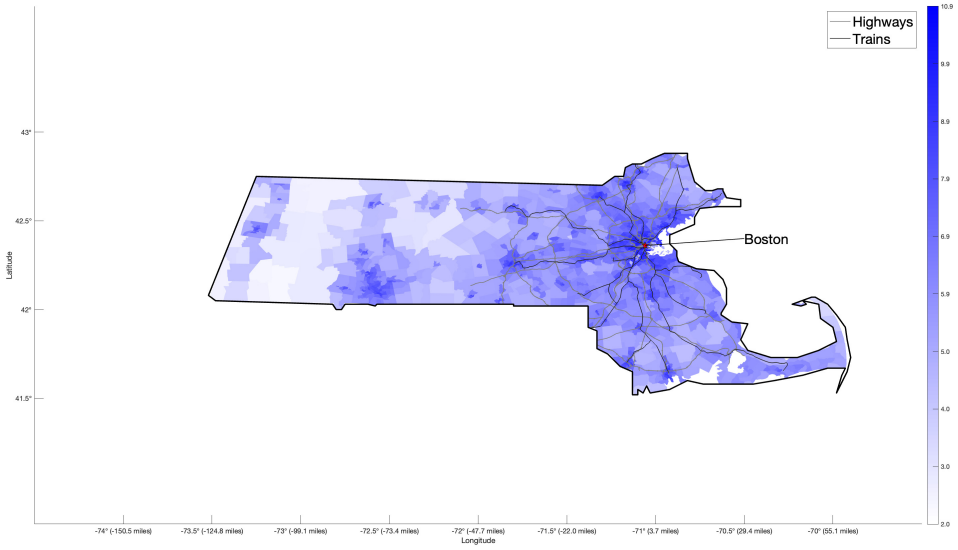
Notes: Estimated driving, location, and public transit usage outcomes in the baseline case without SDVs (Column 1) as well as in counterfactual scenarios when cars are autonomous. Columns 2 through 4 report outcomes when changes in miles driven affect driving speeds. These columns represent different assumptions about the housing supply: Column 2 does not allow any changes in housing, Column 3 allows people to move 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.

Table 4: Impacts of Public Transit Speed Improvements

	Miles Driven		Transit Revenue	
	No SDVs	SDVs	No SDVs	SDVs
Status Quo Transit	0.00%	37.85%	0.00%	-8.12%
	(0.00)	(0.00)	(0.00)	(0.00)
20% Reduction (Improvement)	-1.20%	36.70%	28.91%	17.09%
	(0.00)	(0.00)	(0.00)	(0.00)
40% Reduction (Improvement)	-2.83%	34.98%	69.53%	52.32%
	(0.00)	(0.00)	(0.00)	(0.00)
60% Reduction (Improvement)	-5.36%	32.58%	128.23%	102.78%
	(0.00)	(0.00)	(0.00)	(0.00)
80% Reduction (Improvement)	-8.72%	29.26%	214.13%	176.79%
	(0.00)	(0.00)	(0.00)	(0.00)
100% Reduction (Improvement)	-13.53%	24.27%	342.34%	287.56%
	(0.00)	(0.00)	(0.00)	(0.00)

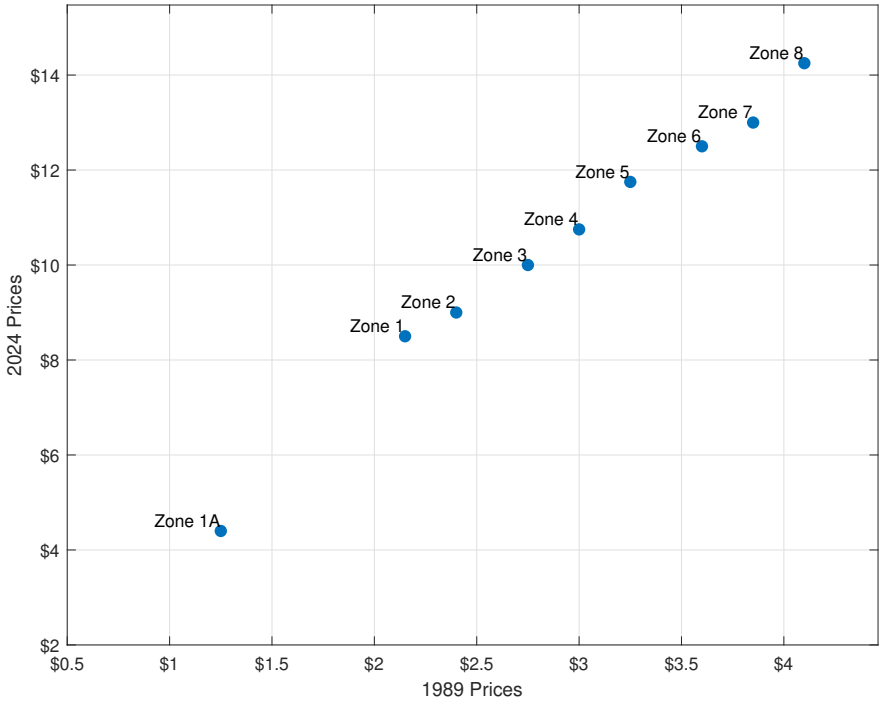
Notes: Simulated changes in miles driven (left two columns) and transit revenue (right two columns) for varying percent improvements in the speed of transit trains, relative to status quo transit speeds with no SDVs. Columns 1 and 3 report outcomes when there are no SDVs, and Columns 2 and 4 show outcomes, relative to the status quo, in the presence of SDVs. Standard errors in parentheses.

Figure 1: Population Density by Census Tract



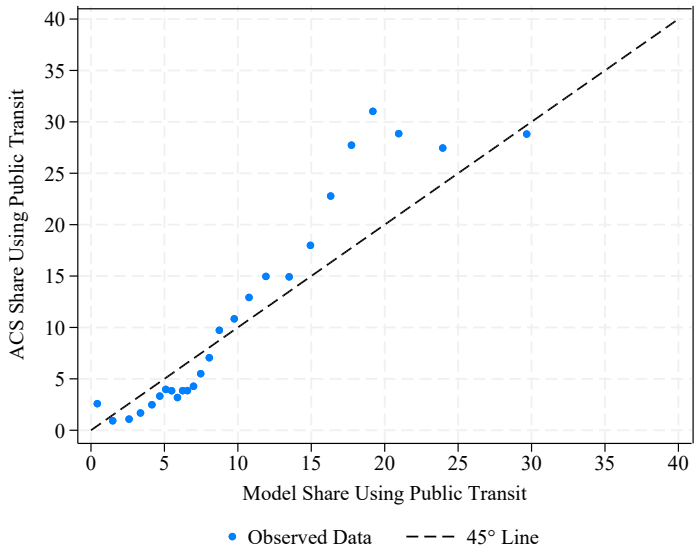
Notes: The map shows the natural log of the home-location population density in workers per square mile of land (excluding bodies of water), by census tract in Massachusetts. Darker shades of blue correspond to more dense census tracts.

Figure 2: Commuter Rail Zone Prices: 1989 vs 2024



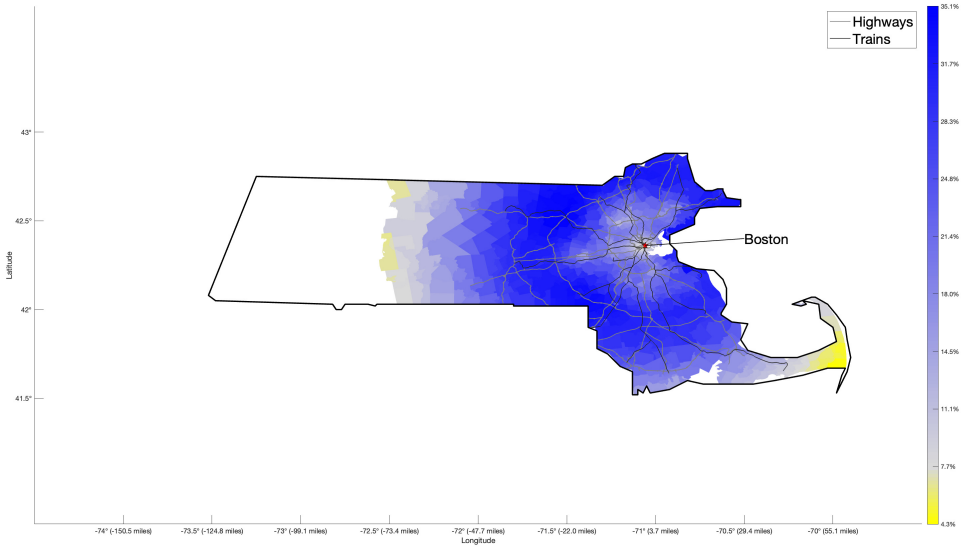
Notes: Nominal commuter rail ticket prices for the 9 commuter rail zones as defined by the Massachusetts Bay Transportation Authority (MBTA), in 1989 (x-axis) vs. in 2024 (y-axis). 1989 prices are taken from the Commonwealth Magazine (2016). See https://commonwealthbeacon.org/uncategorized/pulling-the-data-together/.

Figure 3: Comparison of Model Predictions and ACS Survey Data



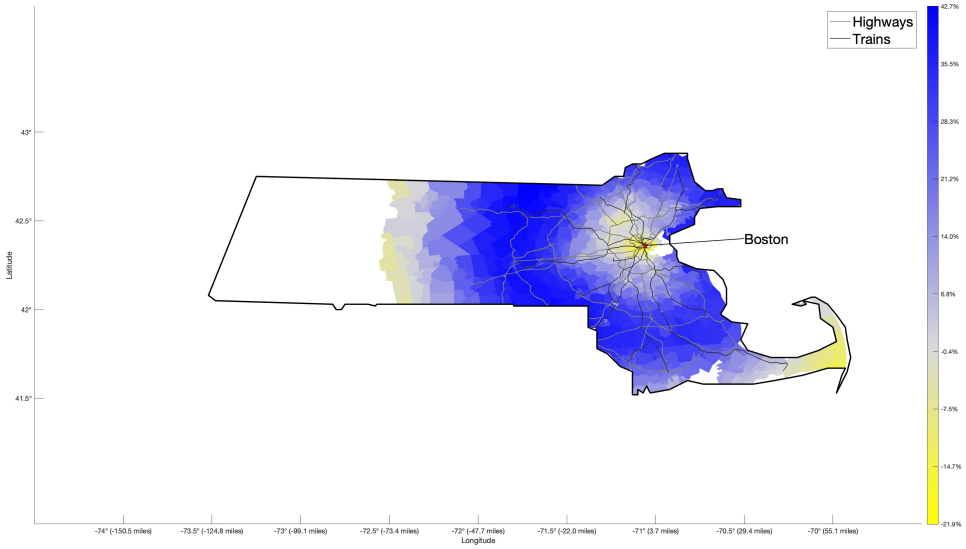
Notes: This figure compares the model's predicted census tract-level share of workers using public transit (x-axis) with 5-year estimates based on survey data from the 2014-2018 American Community Survey (ACS) (U.S. Census Bureau, 2018). Each point represents 50 combined census tracts, grouped based 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 4: Changes in Housing Prices from SDVs with Inelastic Supply



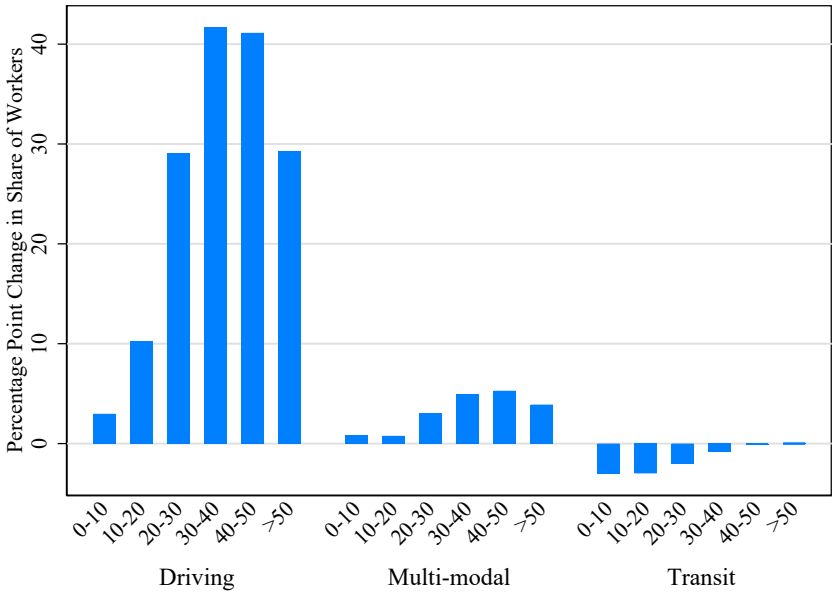
Notes: The map displays, for each census tract, the percentage change in housing prices required to balance demand with the existing housing supply when SDVs are introduced (as in Column 3 of Table 3). Blue shades indicate larger price increases, while yellow shades represent smaller price changes. On average, prices increase everywhere to hold the overall appeal of each location constant despite improved transportation options.

Figure 5: Changes in Worker Home Locations from SDVs with Elastic 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.

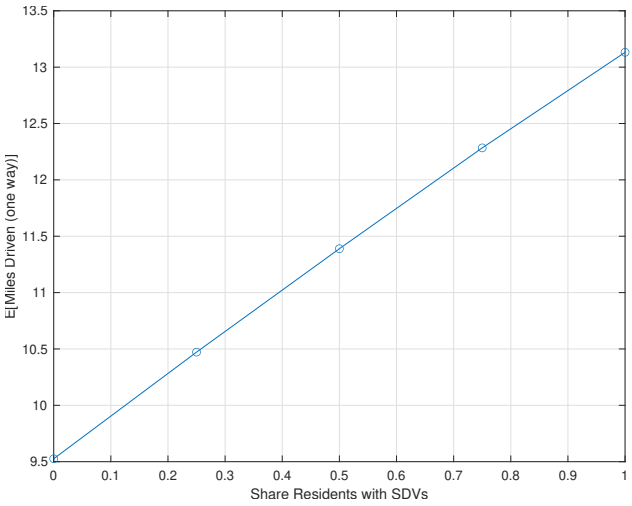
Figure 6: Changes in Transportation Mode Choices from SDVs with Inelastic Supply



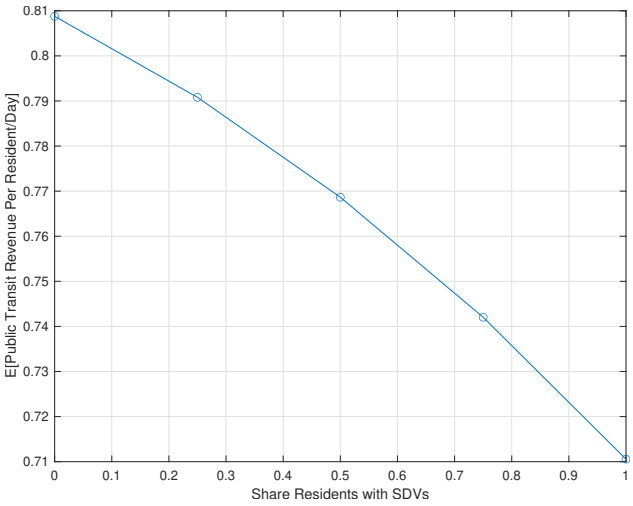
Notes: The graph shows changes in the shares of workers using each of three modes of transportation when SDVs are introduced and home locations are persistent, for different distance ranges, from 0 to 10 miles, to over 50 miles. For example, among people living between 20 and 30 miles from work, the share driving to work rises by 34 points, the share using both cars and transit rises by about 4 points, and the share using only public transit decreases by about 3 points.

Figure 7: Partial Self-Driving Penetration

Panel A: Mileage



Panel B: Transit Revenue



Note: Estimated miles driven (Panel A) and transit revenue (Panel B) for varying levels of SDVs adoption, assuming perfectly elastic housing supply. For each panel, the leftmost point depicts the status quo, with no SDVs (Column 1 of Table 3, and the rightmost point coincides with the outcomes in Column 4 of that table. Points between those two describe scenarios in which the respective shares of cars are autonomous.

A Appendix

A.1 Gradient

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

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

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:

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

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:

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

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

$$\frac{\partial \mathbb{s}_h}{\partial \delta_h} = \int_w \mathbb{s}_{hw} \times (1 - \mathbb{s}_{hw}), \tag{13}$$

$$\frac{\partial \mathbb{s}_h}{\partial \delta_{h' \neq h}} = \int_w -\mathbb{s}_{hw} \times \mathbb{s}_{h'w}, \tag{14}$$

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

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

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 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 1.5 in our main model and 1.25 in this alternative specification. Both ratios confirm commuters' preference for passive transit over active driving, though the magnitude differs slightly between specifications.

Despite these similarities, our main model with the outside transportation option demonstrates superior fit with auxiliary data. Figure A.1 compares the predicted share of public transit users according to both model specifications against corresponding shares from auxiliary data.[42] The predictions from our main model (see Figure 3) align considerably better with observed commuting patterns, 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 very inelastic. When we assume home location is persistent, changes in miles driven by are muted and directly displace transit ridership. See figure Figure A.2.

[42] For more details on the construction of these figures, see Section ??

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 revenue effects regardless of SDV availability, but such improvements cannot fully mitigate the increased vehicle miles traveled resulting from SDV adoption.

For prediction comparison at census tract level, see figure A.1

A.3 Congestion Zone Distances

We employ a great-circle path methodology to estimate the distance traveled within each concentric circle. 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 $M_{hw,1}$ 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:

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

where $D_{hw,c}$ represents the distance within circle c exclusive of inner circles, $M_{hw,c}$ 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 $M_{hw,c}$ 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 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:

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

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 $M_{hw,c}$ 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 circle is zero). Figure A.3 illustrates various scenarios. Approximated driving miles exclusive to each circle is then calculated according to Equation 16.

A.4 Appendix Tables and Figures

Table A.1: Model Estimates – No Outside Option

	Households Earning:	
	$\leq$ 3,333 per month	$>$ \$3,333 per month
1(Involves Public Transit)	0.325	-0.521
	(0.021)	(0.018)
Transfer Count	-0.045	-0.083
	(0.014)	(0.010)
Driving Duration	-0.060	-0.054
	(0.000)	(0.000)
Transit Duration (Non-Auto)	-0.048	-0.043
	(0.001)	(0.000)
Price (\$)	-0.027	-0.037
	(0.001)	(0.001)

Notes: 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: Comparison: Status Quo vs. SDVs Without Outside Options

		Self-Driving Vehicles (SDVs)			
		Dynamic Congestion Effects			Congestion
	Status Quo	Persistent Location Choices	Fixed Housing Supply	Flexible Housing Supply	Same as Status Quo
Transportation Outcomes					
E[Miles Driven (one way)]	15.39	15.60	16.18	19.43	22.91
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
E[Distance From Work]	16.31	16.31	16.91	20.18	23.55
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
E[Public Transit Revenue Per Worker/Day]	1.26	1.16	1.17	1.20	1.16
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
E[Public Transit Revenue Per Rider/Day]	6.14	6.02	6.06	6.08	5.99
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Pct Take Public Transit	20.48	19.28	19.29	19.70	19.40
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$\frac{\text{Speed with Congestion}}{\text{Free-Flow Speed}}$					
Inner Ring	0.57	0.55	0.53	0.52	0.57
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Middle Ring	0.62	0.62	0.61	0.55	0.62
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Outer Ring	0.90	0.90	0.90	0.84	0.90

Notes: 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 does not allow any changes in housing, Column 3 allows people to move 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. Standard errors in parentheses.

Table A.3: Impacts of Public Transit Improvements Without Outside Options

	Miles Driven		Transit Revenue	
	No SDVs	SDVs	No SDVs	SDVs
Status Quo Transit	0.00%	29.87%	0.00%	-4.69%
	(NaN)	(NaN)	(NaN)	(NaN)
20% Reduction (Improvement)	-0.95%	28.60%	24.08%	17.88%
	(NaN)	(NaN)	(NaN)	(NaN)
40% Reduction (Improvement)	-2.47%	27.14%	55.93%	47.68%
	(NaN)	(NaN)	(NaN)	(NaN)
60% Reduction (Improvement)	-4.31%	24.82%	98.58%	87.73%
	(NaN)	(NaN)	(NaN)	(NaN)
80% Reduction (Improvement)	-7.16%	22.13%	156.64%	142.04%
	(NaN)	(NaN)	(NaN)	(NaN)
100% Reduction (Improvement)	-10.86%	17.99%	236.05%	216.59%
	(NaN)	(NaN)	(NaN)	(NaN)

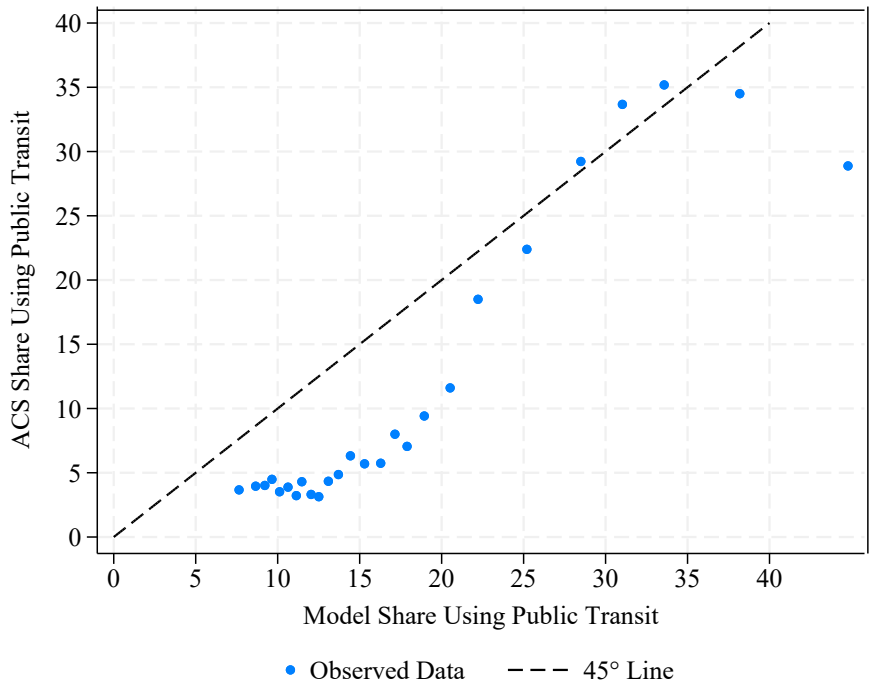
Note: Simulated changes in miles driven (left two columns) and transit revenue (right two columns) for varying percent improvements in the speed of transit trains, relative to status quo transit speeds with no SDVs. Columns 1 and 3 report outcomes when there are no SDVs, and Columns 2 and 4 show outcomes, relative to the status quo, in the presence of SDVs. Standard errors in parentheses.

Table A.4: Impacts of Public Transit Stops Improvements

	Miles Driven		Transit Revenue	
	No SDVs	SDVs	No SDVs	SDVs
Status Quo Transit	0.00%	37.85%	0.00%	-8.12%
	(0.00)	(0.00)	(0.00)	(0.00)
20% Reduction (Improvement)	-0.53%	37.21%	14.47%	4.85%
	(0.00)	(0.00)	(0.00)	(0.00)
40% Reduction (Improvement)	-1.38%	36.51%	32.20%	20.58%
	(0.00)	(0.00)	(0.00)	(0.00)
60% Reduction (Improvement)	-2.22%	35.57%	54.00%	39.98%
	(0.00)	(0.00)	(0.00)	(0.00)
80% Reduction (Improvement)	-3.48%	34.46%	81.65%	64.36%
	(0.00)	(0.00)	(0.00)	(0.00)
100% Reduction (Improvement)	-4.84%	33.13%	117.33%	95.80%
	(0.00)	(0.00)	(0.00)	(0.00)

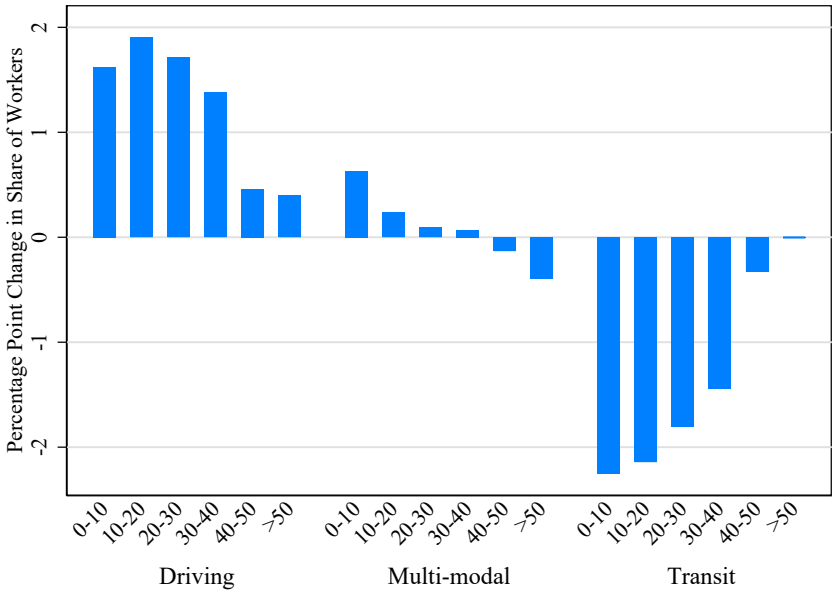
Notes: Simulated changes in miles driven (left two columns) and transit revenue (right two columns) for varying percent improvements in the hassle cost of changing trains, relative to status quo transit speeds with no SDVs. Columns 1 and 3 report outcomes when there are no SDVs, and Columns 2 and 4 show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

Figure A.1: Comparison of Model Predictions and ACS Survey Data – No Outside Option



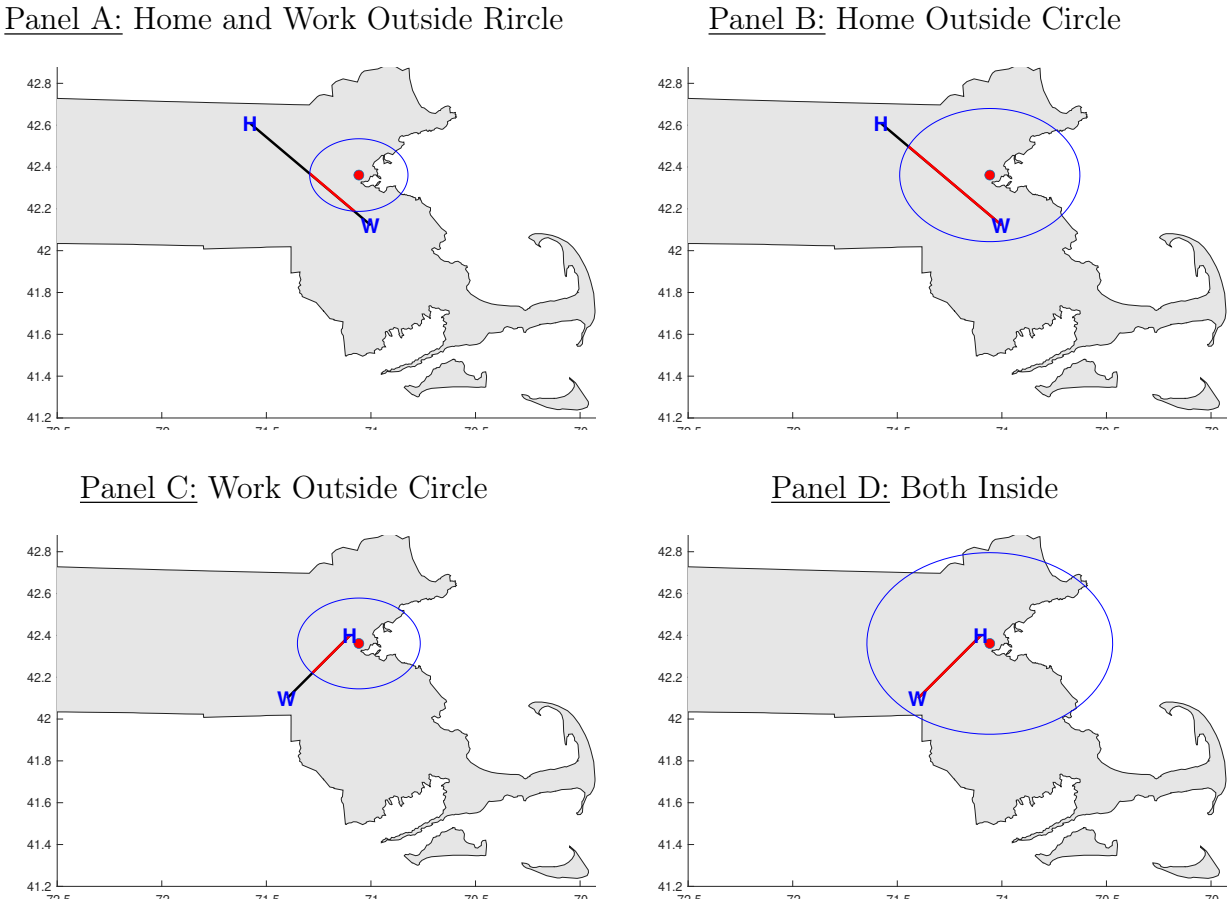
Notes: This figure compares the model's predicted census tract-level share of workers using public transit (x-axis) with 5-year estimates based on survey data from the 2014-2018 American Community Survey (ACS) (U.S. Census Bureau, 2018), when not allowing for an outside option. Each point represents 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: Changes in Transportation Mode Choices from SDVs with Inelastic Supply



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: Depicting Line Portion Within Circle



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. These two points are connected by a black line, representing the worker's commute route. 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.