

Measuring and Mitigating Traffic Externalities*

Zane Kashner[†]

Brad Ross[‡]

(Job Market Paper)

This version: November 10, 2024

[Click here to access the latest version](#)

Abstract

Denser housing construction can alleviate rising housing costs, but opponents frequently cite car traffic as a primary concern. We quantify these heterogeneous traffic costs from new residents across the Boston Combined Statistical Area. Using data on households' intra-metro-area travel, the road network, road speeds, and routing decisions, we estimate monthly traffic counts on every street. We find that moving a house from the 25th percentile of the distribution of nearby street traffic to the 75th percentile decreases its value by 20%, while adding the same number of monthly trips to the street of a similar house at the 75th percentile only decreases its value by 2.7%. We estimate a structural, hedonic model of households' residential choices and visits to points of interest and find that households are willing to pay to avoid both car traffic on their street and travel time, but that these preferences vary widely across the population. Using the model and estimates of how traffic volumes affect road speeds, we simulate the traffic externalities caused by adding new residents in different locations. We find that a Massachusetts state law targeting a 10% housing stock increase to land near public transit stops causes \$3.3 billion in traffic externalities from these new residents, an \$820 million reduction relative to spreading those homes uniformly across space. Building those units on thoroughfares instead would decrease welfare costs by an additional \$520 million.

*Authors are listed alphabetically. We thank Brad's advisors Rebecca Diamond, Guido Imbens, and Shoshana Vasserman for their generous support, encouragement, and kindness, and we thank Hunt Allcott, Susan Athey, Lanier Benkard, Liran Einav, Matthew Gentzkow, Neale Mahoney, Peter Reiss, David Ritzwoller, Gregory Rosston, Paulo Somaini, Ali Yurukoglu, and many seminar participants for their insightful comments and valuable feedback. Brad acknowledges and appreciates financial support from the Bradley Graduate and Postgraduate Fellowship via grants to the Stanford Institute for Economic Policy Research. Computational Support was provided by the Data, Analytics, and Research Computing (DARC) group at the Stanford Graduate School of Business ([RRID:SCR_022938](#)).

[†]Stanford Graduate School of Business; Email: zkashner@stanford.edu.

[‡]Stanford Graduate School of Business; Email: bradross@stanford.edu.

1 Introduction

Many policy makers and economists argue that local land use regulations raise housing costs by restricting housing supply, reducing economic growth and equity (Duranton and Puga, 2023; Glaeser and Gyourko, 2018; Hsieh and Moretti, 2019; Mast, 2023). Despite this, many local governments maintain constraints on denser housing construction at the behest of their constituents. A frequently cited concern about new housing developments are their car traffic impacts. For example, 63% of Sacramento Valley residents said an increase in car traffic was their top concern about the construction of multifamily housing in their neighborhood (Valley Vision, 2023), and 75% of California city council meetings about housing or land use also mention car traffic (Martin and Venugopal, n.d.). At the same time, states like Colorado and Massachusetts have successfully enacted laws requiring municipalities to plan for targeted denser housing near public transit stops.¹ Quantifying the heterogeneous traffic costs of new housing units can allow policy makers to make planning and siting decisions in part to minimize them.

In this paper, we develop a methodology for measuring spatially heterogeneous costs of traffic exposure and congestion to households. These estimates allow us to quantify the traffic externalities imposed on incumbent households by adding new drivers in different locations. We infer costs to residents by building and estimating a structural, hedonic model of households' location choices and visits to points of interest (POIs) like offices, schools, and consumption amenities in the Boston Combined Statistical Area (CSA). We estimate households' heterogeneous willingness to pay to avoid cars driving down their streets and to avoid spending more time traveling to POIs using the causal effects on housing transaction prices of increased car trips on their streets and intra-metro-area travel times.

With estimates of preferences and the effects of increasing traffic volume on equilibrium road speeds, we simulate the heterogeneous welfare costs to incumbent residents from adding a new resident and their associated car traffic to each street across the Boston CSA. We then evaluate the traffic externalities caused by variants of Massachusetts' 2021 "MBTA Communities" law, which required municipalities across eastern Massachusetts to plan for a cumulative 10% increase in metropolitan area population near public transit stops. We find that siting as specified in the law saves around \$820 million in traffic externalities relative to uniformly allocating that new housing across space. Concentrating this increased density on thoroughfares instead would further reduce traffic externalities by approximately \$580 million.

¹Massachusetts passed its "MBTA Communities" law in January, 2021, and Colorado passed its "Housing in Transit-Oriented Communities" law in May, 2024.

To estimate households’ desires to avoid traffic on their streets and the effect of increasing traffic volume on travel times, we first construct estimates of counts of monthly car trips through every street in the Boston CSA in 2010 and 2018. Most existing datasets only provide data on traffic counts for subsets of larger, non-residential roads, making it difficult to assess the impacts of increased traffic flow on residential streets where most households live. To construct our traffic counts, we first estimate monthly car traffic flows between every house and every POI in the Boston CSA using GPS data on which POIs different households visit along with data on the universe of residential property locations and travel mode choices. We then combine a snapshot of the Boston CSA road network with a large sample of roads’ speed distributions in 2010 and 2018 and compute the billions of optimal car routes all households take when making POI visits across different draws from roads’ speed distributions. Finally, we sum up these per-household trip counts along their optimal routes to construct our street-level traffic counts.

Next, we characterize households’ distastes for exposure to car traffic on their streets. We take a hedonic approach and estimate the causal effect of an increase in monthly car trips down a street on the transaction prices of houses on that street. A key challenge is that unobserved changes in home and neighborhood quality could be correlated with changes in nearby traffic and POI accessibility over time.² The direction of this bias is theoretically ambiguous; for example, increases in traffic over time could be driven by factors that raise house prices, like improved nearby amenities, or factors that lower house prices, like increased supply of housing nearby.

To circumvent these issues, we construct a shift-share instrument that takes advantage of differences in the substitutability of households’ optimal routes across metropolitan-area-wide traffic states. For intuition, suppose there are two traffic states, “low” and “high,” that each occur with some probability. For some households, alternative routes exist in the high traffic state that allow them to avoid most of the high-traffic slowdown along their low-traffic optimal route. For others, no alternative routes exist. As such, changes in the probabilities of these traffic states between 2010 and 2018 due to unforeseen shocks like the advent of ride-sharing and smartphone-based routing apps will induce differential changes in traffic counts across streets, even if households’ origins and destinations are held fixed. Since we construct our traffic counts in 2010 and 2018 by taking sums of state-specific estimated traffic count estimates weighted by estimates of these traffic state probabilities in each year, we can instrument for changes in car traffic with the changes that occur from holding state-specific traffic flows fixed in 2010 and only allowing the traffic state probabilities to vary.

²This relationship is potentially more severely confounded in the cross section, where housing units on busier streets could be of lower quality in unobserved ways.

To accommodate potential nonlinearities in the effect of traffic on houses’ streets on their transaction prices, we use a nonparametric instrumental variables estimator (X. Chen and Christensen, 2018; Newey and Powell, 2003). We find that moving the same house from the 25th percentile of monthly trips down a house’s street to the 75th percentile causes a 20% decrease in home value, while adding the same quantity of trips to the street of a house at the 75th percentile of nearby monthly trips only decreases that house’s value by 2.7%. To maximize power, we control for housing unit characteristics via random forests in a partially linear fashion (Athey, Tibshirani, and Wager, 2019; Chernozhukov et al., 2018), and we control for street and Census block group-by-year unobservables using generalizations of fixed effects introduced in Arkhangelsky and Imbens (2022) and Arkhangelsky and Imbens (2023).

Next, we translate these causal effect estimates into estimates of households’ willingnesses to pay to avoid additional monthly trips down their streets. To do so, we specify a flexible, hedonic model of households’ housing choices in which households derive heterogeneous utilities from each housing unit’s nearby traffic exposure along with other property and neighborhood characteristics (Bajari and Benkard, 2005; Berry and Pakes, 2007; Gorman, 1980; Lancaster, 1966; Rosen, 1974). Given the large number of housing options in the Boston CSA, we view households as making approximately continuous choices over bundles of housing characteristics. This perspective allows us to apply seminal results from Bajari and Benkard (2005) and use the derivatives of our nonparametric house price surface estimate with respect to monthly street traffic as our estimates of households’ disutilities from nearby traffic exposure. We find that the median household is willing to pay around \$3.50 to avoid an additional monthly car trip on their street, but that willingness to pay can be larger or smaller by orders of magnitude across the population, with the 10th and 90th percentiles of households willing to pay \$0.05 and \$32 per avoided nearby monthly car trip.

We also estimate households’ heterogeneous distastes for travel time to POIs of different types. While many papers have estimated individuals’ value of time using short-run variation in travel times from fixed locations,³ our approach allows us to estimate households’ heterogeneous values of time for visits to different types of POIs from their long-run choices of neighborhoods with differential access to POIs.⁴ We develop a model of households’ monthly decisions to visit POIs across the Boston CSA. First, households with heterogeneous preferences over visits to different categories of POIs choose how much time to spend visiting POIs in each category. Then, they choose how to allocate that time to visiting POIs within

³See Section 1.1 for examples.

⁴As discussed in detail in Section 1.1, many papers also estimate households’ value of time from their residential and/or commuting choices. Our hedonic approach enables us to estimate richer value-of-time heterogeneity across households.

each category across the Boston CSA, where POIs compete for visitors by provisioning visit quality in a monopolistically competitive fashion. Importantly, travel times are random and depend on the frequencies of different traffic states. Under this model, a household’s indirect utility for POI visits from a particular home neighborhood can be written as a household-specific weighted sum over home-neighborhood and category-specific “travel time indices,” which characterize the visit-quality-adjusted difficulty of accessing POIs in that category for any household that chooses to live in that neighborhood. Because the travel time indices for each neighborhood are common across households and only the weights households place on them are household-specific, we can nest this POI choice problem in our hedonic utility framework, conceptualizing a household as valuing access to POIs from a home neighborhood’s hedonically via its low-dimensional bundle of travel time indices.

We use our POI visit count data from SafeGraph to recover the parameters governing households’ choices of visits to POIs. We show that the model admits a reduced form for the number of visits per POI in a category and neighborhood that enables estimation via an algorithm akin to a Poisson regression with high-dimensional fixed effects (Bergé, 2018), but with a mixture term capturing expected travel time disutility with respect to the distribution of traffic states. We assume households living in the same Census block group have the same preferences over POIs, which allows us to recover households’ preferences over POI categories from their model-implied, quality-adjusted times spent visiting POIs in each category. Finally, we translate these POI visit utilities into hedonic willingness to pay by estimating the causal effects of increases in a neighborhood’s travel time indices on house prices in that neighborhood in a similar nonparametric fashion to our approach for estimating car traffic exposure disutilities. Because changes in travel times to POIs could be correlated with unobserved changes in neighborhood quality, we again instrument for changes in our travel time indices with changes in the indices induced only by changes in traffic state probabilities. We find that households’ values of time implied by our estimates vary substantially in both absolute terms (the 10th and 90th percentiles are around \$12/hour and \$199/hour) and relative to median hourly wages in their neighborhood (the 10th and 90th percentiles are 27% and 511% of neighborhood median hourly wages).⁵

The last ingredient required to simulate the congestion externalities from new residents is the effect of increased traffic volume on road speeds and, by extension, travel times. We estimate these congestion functions using changes in road speeds and traffic densities between 2010 and 2018. Since traffic speeds and densities are determined simultaneously

⁵These estimates are broadly in line with other value of time estimates based on households’ joint residential and commuting choices (see Section 1.1 for references). See Su (2022a) for one summary of the constellation of approaches to and estimates of values of travel time.

in equilibrium and drivers decide whether to take a trip and which route to take based on prevailing traffic conditions (Akbar, Couture, Duranton, and Storeygard, 2023; Couture, Duranton, and Turner, 2018; Hall, 1996; Small and Chu, 2003; Yang, Purevjav, and Li, 2020), we instrument for changes in traffic densities with changes just induced by differences in aggregate traffic state probabilities over time, as before. After converting these effects of traffic densities on speeds into effects of traffic volume on travel times as in Yang et al. (2020), we find that most street-level traversal times are relatively unresponsive to changes in traffic volume, albeit more responsive than Yang et al. (2020).

Having estimated households' disutilities from exposure to car traffic on their street and from increased travel times due to increased congestion, we simulate the traffic externalities caused by a counterfactual new resident added to each street in the Boston CSA. Adding a new resident to a street who drives does not just affect the other households living on that street; any households living on the streets on which that new resident drives will be exposed to their car trips, and any households driving on the same streets will also be slowed down by congestion caused by the new resident's car trips. We assume this simulated resident visits the same POIs and travels by car at the same frequency as the other residents on the same street. Adding an additional resident to the median plot of residential land in the Boston CSA imposes around \$11,100 in traffic externalities on incumbent households, but this cost varies substantially across space. At the 10th percentile, a new resident causes only around \$5,600 in traffic externalities, while at the 90th percentile, a new resident causes nearly \$21,800 in traffic externalities. Exposure costs account for a substantial share of these externalities; for half of potential locations, the welfare costs from exposure externalities are greater than costs from congestion externalities, and the distribution's right tail is large.

Finally, we use these estimates to evaluate the traffic externalities caused by new housing built under variants of Massachusetts' 2021 MBTA Communities law. Every municipality served by the MBTA, except Boston, was required by law to change its land use regulations to allow for a significant increase in population within half of a mile of public transit stations, amounting to a 10% planned increase in the metropolitan area's housing stock. This targeted upzoning policy is different than proposals in other American cities and states that mandate more uniform loosening of building restrictions across space.⁶ We compare the traffic impacts of adding each municipality's allotted housing units according to three coarse targeting policies: within a half mile of public transit stops as the original law dictates, across all residential land within its borders, or instead along larger streets.

We find that building new housing in the locations specified by the original MBTA

⁶For example, the city council in Cambridge, Massachusetts recently proposed allowing six-story apartment buildings on every residential lot within their borders (Cambridge City Council, 2024).

Communities law would cause a \$3.3 billion welfare loss to Boston CSA incumbents from traffic externalities, \$2.1 billion (64%) of which are exposure externalities. In contrast, building the same quantities of new housing in each municipality but spreading them across all residential land more evenly would cause \$4.1 billion in traffic externalities, \$2.6 billion (64%) of which are exposure externalities. Around \$500 million of the \$820 million in additional traffic externalities (61%) are due to increases in exposure externalities. While targeting new housing near public transit meaningfully reduces traffic externalities relative to increasing density across residential lots, its distributional impacts are mixed. Under the MBTA Communities targeting policy, traffic externalities are 22% larger in block groups with bottom quartile median household incomes, while traffic externalities are 27% smaller in the top quartile of block groups by income. If the MBTA Communities law’s housing units were built instead on residential land on thoroughfares, incumbents would incur only \$2.7 billion in traffic externalities, with only 48% driven by exposure externalities. While this targeting rule only reduces congestion externalities by \$50 million relative to the policy variant with minimum density increases everywhere, households in every neighborhood income quartile are hurt less by this policy variant’s traffic externalities than in the other two we consider.

1.1 Related literatures

This paper contributes to several literatures in urban and public economics. Most directly, a body of existing work has estimated various costs of exposure to car traffic like pollution (Currie and Walker, 2011), noise (Svärdh and Genell, 2020), lack of safety (Ewing and Dumbaugh, 2009), and parking trouble (Shoup, 2018). This paper remains agnostic about the source of households’ distaste for traffic exposure and instead estimates households’ hedonic disutilities from the entire bundle of costs nearby traffic exposure can impose. Recently, Tang (2021) also estimated substantial hedonic value of decreased traffic exposure; we are able to characterize heterogeneity in these costs across households and study its implications for counterfactual planning policies.

In addition, a large body of work has quantified the costs of car traffic congestion and policies intended to reduce it via infrastructure improvements and or road pricing (Allen and Arkolakis, 2022; Almagro, Barbieri, Castillo, Hickok, and Salz, 2024; Ater, Shany, Ross, Turkel, and Vasserman, n.d.; Barwick, Li, Waxman, Wu, and Xia, 2024; Bordeu, 2023; Cook and Li, 2023; Durrmeyer and Martinez, 2022; Herzog, 2024a, 2024b; Kreindler, 2024; Severen, 2023; Tsivanidis, 2024; Yang et al., 2020). An important determinant of many of these traffic cost estimates is peoples’ values of travel time, which a rich literature also seeks to estimate accurately via both “short-run” variation in traffic conditions (Buchholz, Doval,

Kastl, Matějka, and Salz, 2020; Castillo, 2023; Goldszmidt et al., 2020; Rosaia, n.d.; Small, Winston, and Yan, 2005) and “longer-run” choices of trip destinations and or origins (Su, 2022b). We estimate values of time that are broadly consistent with the latter literature.

Our work also relates to the literature on spillovers from increased housing density. Recent work has found that increasing (allowable) density decreases nearby rents and property values (Anagol, Ferreira, and Rexer, 2021; Asquith, Mast, and Reed, 2023; Davidoff, Pavlov, and Somerville, 2022; Diamond and McQuade, 2019; Li, 2022; Pennington, 2021; Severen and Plantinga, 2018; Turner, Haughwout, and Van Der Klaauw, 2014). However, these papers provide mixed evidence as to whether these price decreases are due primarily to the “supply effect” of increasing the quantity of similar housing units or an “externality effect” of adding density on neighborhood character and/or nearby amenities. Pollmann (2020) shows that using variation in distance to the same density shock for identification as many of these papers do can end up differencing out much of the effect of the shock if its impact diffuses broadly in space, as car traffic tends to do. Instead, we measure traffic costs directly and use traffic simulations to characterize the spatial extent of a particular negative externality from increased housing density.

More broadly, our work is situated in a large literature that characterizes how much households value certain housing and neighborhood characteristics via the effects those characteristics have on housing values, e.g. Banzhaf (2021); Campbell, Giglio, and Pathak (2011); Chay and Greenstone (2005); Diamond and McQuade (2019); Greenstone and Gallagher (2008); Han, Hebllich, Timmins, and Zylberberg (2024); Linden and Rockoff (2008); Rossi-Hansberg, Sarte, and Owens III (2010). We use state-of-the-art causal inference methods to identify households’ heterogeneous distastes for traffic exposure and congestion from the curvature of an estimated hedonic price surface (Agarwal, Li, and Somaini, 2023; Bajari and Benkard, 2005; Banzhaf, 2021; Rosen, 1974).

Finally, a large recent literature has developed integrated models of residential choice and intra-metro area travel (Ahlfeldt, Redding, Sturm, and Wolf, 2015; Almagro and Dominguez-Iino, 2022; Athey, Blei, Donnelly, Ruiz, and Schmidt, 2021; Athey, Ferguson, Gentzkow, and Schmidt, 2021; Barwick et al., 2024; Bordeu, 2023; Cao, Chevalier, Williams, and Parsley, n.d.; Couture, Gaubert, Handbury, and Hurst, 2023; Dingel and Tintelnot, 2020; Duranton and Puga, 2023; Miyauchi, Nakajima, and Redding, 2021; Monte, Redding, and Rossi-Hansberg, 2018; Oh and Seo, 2022; Severen, 2023; Tsivanidis, 2024). Given our counterfactuals of interest, our hedonic approach enables estimation of rich heterogeneity in households’ preferences for both traffic exposure and access to POIs of all types without needing to specify households’ preferences for other determinants of housing and neighborhood desirability. Our model of households’ allocation of time to visit different types of POIs most resembles

models proposed by [Couture \(2016\)](#), [Su \(2022a\)](#), and [Cook \(2022\)](#) and our monopolistically competitive structure resembles the amenity supply model in [Almagro and Dominguez-lino \(2022\)](#), but we aim to capture additional heterogeneity across households in their category preferences.

2 Data and measurement

To quantify households’ disutility from exposure to car traffic and to estimate how traffic volumes affect road speeds, we first need measures of traffic flows down streets and a notion of households’ exposures to those flows. Such data are hard to come by at scale,⁷ particularly on the residential streets on which most households live. As such, we construct our own measure by aggregating several sources of data on traffic flows and conditions in the Boston-Worcester-Providence Combined Statistical Area (CSA), a group of counties defined by the United States Office of Management and Budget that represents a unified labor market encompassing eastern Massachusetts, all of Rhode Island, southern New Hampshire, and northeast Connecticut.

2.1 Data sources

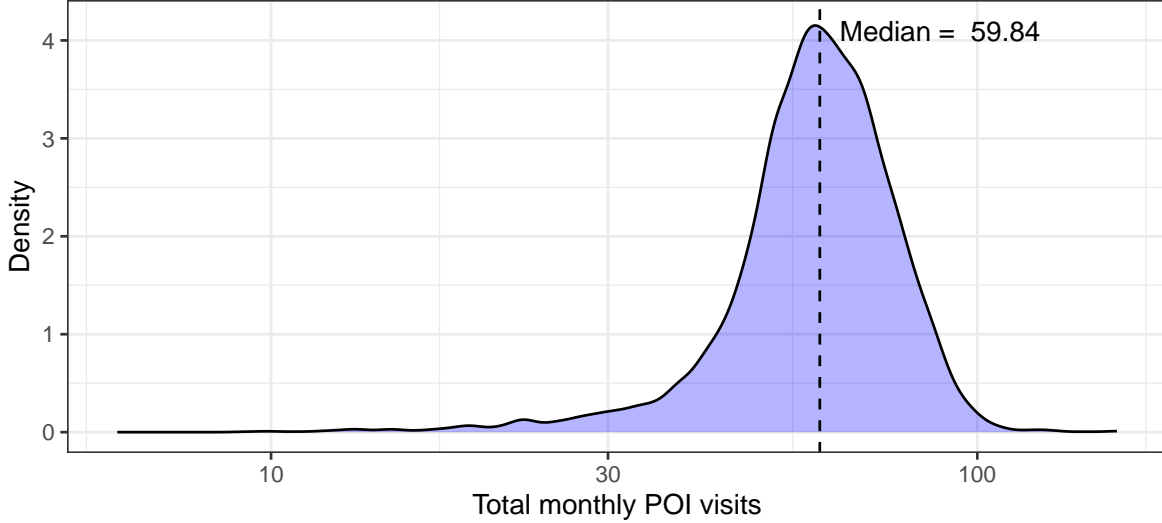
SafeGraph Places and Visits. SafeGraph provides a dataset based on GPS pings from a nationally-representative sample of smartphones containing monthly counts of visits to particular points of interest (POIs) by devices owned by people that reside in different 2010 Census block groups.⁸ SafeGraph also provides POIs’ geographic coordinates, NAICS code classifications, and median time spent “dwelling” at them. A benefit of measuring origin/destination flows using these data is that they are more representative of all POI visits, not just the 15% of trips that are commutes ([USDOT BTS, 2017](#)).⁹ We use a cross-section of trip counts from April of 2018. [Figure 1](#) presents the distribution of number of trips in the Boston CSA. For reference, the median person in the data takes 60 trips per month, and people frequently travel several miles from home to visit even frequently-visited POIs like grocery stores and schools. When we estimate households’ preferences for access to POIs in different categories, we group POIs into seven categories based on their two-digit

⁷Most cities only conduct traffic studies infrequently and on main thoroughfares.

⁸Census block groups are constructed to include between 600 and 3,000 people.

⁹In particular, the 2017 National Household Transportation Survey finds that 15% of daily trips are commuting trips, 45% of daily trips are for shopping and/or errands, and 27% of daily trips are social in nature ([USDOT BTS, 2017](#)).

Figure 1: The distribution of trip frequencies



NAICS codes.¹⁰ Figure A.1 shows the distributions across devices of counts of visits to POIs in these seven categories for people in the Boston CSA.

CoreLogic Tax and Deed Datasets. CoreLogic’s historical tax dataset contains a myriad of characteristics of every property in the United States. In particular, we observe each property’s geographic coordinates, number of housing units on the property, owner occupancy,¹¹ and observable characteristics like lot size, building age, square footage per unit, and counts of bedrooms, and bathrooms, and other rooms per unit. CoreLogic’s ownership transfer dataset also contains data on the universe of property transactions since 1997, including transaction prices for all states with mandatory transaction price reporting, as Massachusetts, New Hampshire, Rhode Island, and Connecticut all do.

US Census and American Community Survey. We derive counts of adult drivers by Census block group in 2010 and 2018 using data from the 2010 Census and 2014-2018 American Community Survey (ACS) aggregated tables. Since not all multifamily properties in CoreLogic have unit counts listed, we also estimate these unit counts using multifamily building size data by block group, again from the 2014-2018 aggregated tables.

¹⁰The categories grouping two-digit NAICS codes we construct are Industrial: 11, 22, 23, 31, 32, 33; Transportation: 48, 49; Healthcare: 62, 81, 92; Offices: 51 through 56; Education: 61; Stores: 42, 44, 45; and Leisure: 71, 72.

¹¹We infer owner occupancy if the address of the property is the same as the address from which property taxes for that property are paid.

Open Street Map. To determine which streets drivers take when making the POI visits we see in the data from SafeGraph, we use a crowdsourced snapshot of the entire Boston CSA road network from Open Street Map. We also match each property from CoreLogic and POI from SafeGraph to its closest street in the road network.

TomTom Traffic Stats. To know how fast cars actually travel along roads in the Boston CSA, we assemble a dataset from TomTom containing the distributions of cars’ speeds on a large sample of road segments in the Boston CSA over the course of April, 2010 and April, 2018. TomTom constructs such information by aggregating GPS data from their in-car navigation devices, navigation systems licensed to car manufacturers, smartphones, and partner commercial vehicle fleets. For each road segment and time period we sample, we also observe the total number of cars whose data TomTom used to estimate the road segment’s distribution of speeds during that time period.

2.2 Estimating monthly car traffic flows through streets

We now describe how we combine these data sources to construct our measures of monthly traffic flows through streets. First, we construct estimates of monthly car trip counts between every house and POI in the Boston CSA. To do so, we assume that every device residing in a given Census block group takes the same monthly trips to POIs and compute the average number of monthly trips taken to each POI by a device residing in each block group. We then multiply those average monthly trip counts per household for a given Census block group by estimates of the number of drivers residing in each residential property derived from the 2010 Census and 2014-2018 ACS; doing so yields counts of 2010 and 2018 monthly trips taken from every residential property to every POI in the Boston CSA.¹²

Next, we infer which roads households take when making these car trips. To do so, we repeatedly choose sets of speeds for all roads from the road speed distributions provided by TomTom in a manner described in detail in Appendix B. We call each profile of speeds across the Boston CSA’s roads a *traffic state*. For each of these traffic states, we compute the optimal driving route taken by every household to every POI they visit, amounting to tens of billions of routes. We then add up the total monthly car traffic flows through each street in each traffic state to construct traffic-state-specific counts of monthly cars passing through each street. To aggregate these state-specific traffic counts into monthly total street-level car traffic counts, we sum these state-specific flows multiplied by the probabilities of

¹²Since we only have POI visiting data for 2018, these origin/destination flows will still reflect POI visiting patterns so long as the locations of POIs in space are similar in 2010 and 2018, even if the exact POIs are not.

being in each state; see Appendix B for details on how we construct those probabilities. We repeat this procedure using our state-specific flows and speeds in 2010 and 2018 to construct separate estimated monthly traffic counts for each Boston CSA street in 2010 and 2018. In Figure A.2, we plot the 2018 distribution of our inferred monthly traffic counts across the nearest streets of each housing unit in the Boston CSA.

3 Valuing households’ traffic exposure costs

To assess how much households care about avoiding cars driving past their homes, we estimate the causal effect of increasing the number of monthly car trips through a street on the prices of homes on that street. We do so by relating changes in street-level car trips between 2010 and 2018 to changes in the transaction prices of homes on those streets over the same time period. We use a shift-share instrument to isolate exogenous variation in these street traffic changes due just to differences in streets’ exposures to changes in the frequency of aggregate traffic states, as we describe in more detail in Section 3.1.1. To accommodate nonlinearities and maximize power, we estimate the effect nonparametrically, and we control flexibly for a high-dimensional set of housing unit characteristics and street and neighborhood-by-year level unobservables using tools at the intersection of causal inference and machine learning. We then interpret these causal effects through the lens of a hedonic model of households’ long-run residential choices, allowing us to estimate households’ heterogeneous willingness to pay to avoid traffic on their street based on their observed choices.

3.1 Estimating the effect of traffic flows on property values

To estimate the effect of own-street traffic flows on house prices, we use data on 291,147 single-unit residential property transactions from the years 2010-2012 and 2017-2019 from CoreLogic.¹³ We assume our estimates of 2010 traffic flows reflect street traffic conditions for houses transacted in 2010 through 2012 and that our estimates of 2018 traffic flows reflect street traffic conditions for houses transacted in 2017 through 2019. We leave out transactions between 2013 and 2016 because we cannot estimate traffic flows that would accurately reflect traffic conditions those years given that we only have cross-sections of road

¹³We restrict our attention to arms-length transactions of owner-occupied single-family homes and single-unit condominiums since prices of rental properties may be set differently and unit counts for multi-unit apartment buildings in CoreLogic can be unreliable (albeit less so in New England than in other regions of the United States).

speed data from 2010 and 2018.¹⁴

To describe our estimation strategy formally, we first introduce some notation. We index housing units by j and years by y , where each housing unit j is associated with a street $\ell(j)$ and a 2010 Census block group $n(j)$, and each year y is associated with a given traffic period $t(y) \in \{2010, 2018\}$, as described above. We let P_{jy} denote the sale price of housing unit j transacted in year y , we let $C_{\ell(j)t(y)}$ denote our estimate of the count of monthly car trips through house j 's street $\ell(j)$ in period $t(y)$, and we let $A_j \in \mathbb{R}^d$ denote a vector of observable housing unit characteristics.¹⁵

We model a housing unit j 's transaction price in year y as being determined by the following flexible function of its observed and unobserved characteristics:

$$P_{jy} = p_{t(y)}^{(C)}(C_{\ell(j)t(y)})p_y^{(-C)}(A_j, \phi_{\ell(j)}, \varphi_{n(j)y}) \exp(\xi_{jy}), \quad (1)$$

where $p_{t(y)}^{(C)}$ is a period-specific, unrestricted function of housing unit j 's exposure to monthly car trips on its street in period $t(y)$,¹⁶ and $p_y^{(-C)}$ is a year-specific, unrestricted function of housing unit j 's observed characteristics, as well as unrestricted street $\ell(j)$ -specific and Census block group $n(j)$ -by-year y unobservables that act as generalizations of street and neighborhood-by-year fixed effects (Arkhangelsky and Imbens, 2022, 2023). ξ_{jy} captures remaining unobserved, within-neighborhood determinants of time-varying housing unit desirability. The model (1) is a nonparametric generalization of a standard log-linear model of house prices (see e.g. Banzhaf (2021)):

$$\log(P_{jy}) = \beta_{t(y)} C_{\ell(j)t(y)} + \pi'_y A_j + \phi_{\ell(j)} + \varphi_{n(j)y} + \xi_{jy},$$

where street traffic $C_{\ell(j)t(y)}$ and housing unit observables affect prices linearly with time-varying coefficients along with street $\ell(j)$ and Census block group $n(j)$ -by-year y fixed effects.

While (1) does restrict heterogeneity in street traffic's effect on transaction prices across houses, we do not require the effect to be linear, and we maintain a healthy amount of flexibility in how variables other than nearby street traffic affect transaction prices. This flexibility enables us to maximize power by stripping away as much extraneous variation in transaction prices as possible. Since not every housing unit is transacted in every year, including time-invariant, street-level "fixed effects" in the model ensures that our estimates

¹⁴We center the first three-year bin around 2011 rather than 2010 to avoid the brunt of the Great Recession's effect on home prices.

¹⁵In particular, A_j contains the age of the housing unit, lot size, floor space, number of bedrooms, number of bathrooms, total number of rooms, and an indicator for whether the housing unit is a single-family home.

¹⁶As discussed in Kuminoff and Pope (2014) and Banzhaf (2021), allowing for time-varying effects is important since, as we will discuss in Section 3.2, not doing so will conflate changes in housing unit characteristics and changes in the equilibrium relationship between house prices and characteristics over time.

of this model will be based on variation from changes in street traffic and transaction prices over time in a repeated cross-sections sense (Banzhaf, 2021), since, within a cross section, it is reasonable to suspect that housing units on busier streets could be of lower quality in unobserved ways.

Despite being able to account for a variety of observed and unobserved differences in housing unit desirability when we go to estimate (1), changes in street traffic over time could still be correlated with unobserved within-neighborhood changes in the desirability of living in houses on that street that could bias estimates of $p_{t(y)}^{(C)}$ in theoretically ambiguous ways. For example, the construction of a new housing development in a neighborhood would cause prices nearby to go down due to the increase in housing supply, but the increase in residents could also increase car traffic on nearby streets. Conversely, if amenities within the neighborhood improved, then a street nearby might get more car traffic from visitors to those amenities and desirability of housing units on that street could improve given their proximity to the improved amenities. To avoid this spurious variation, we construct a shift-share instrument $\tilde{C}_{\ell(j)t(y)}$ for street traffic flow $C_{\ell(j)t(y)}$ that we describe in detail in Section 3.1.1.

Before introducing our instrument, we first describe how we identify and estimate the effect of street car traffic on house prices $p_{t(y)}^{(C)}$ under the model (1). Concretely, we assume that the time-varying house-specific unobservables ξ_{jt} are mean-independent of the instrument $\tilde{C}_{\ell(j)t(y)}$, observable housing unit characteristics A_j , and street and block group-by-year means of transformations of these variables $\bar{X}_{\ell(j)}$ and $\bar{X}_{n(j)y}$, respectively:

$$\mathbb{E} \left[\xi_{jt} \mid \tilde{C}_{\ell(j)t(y)}, A_j, \bar{X}_{\ell(j)}, \bar{X}_{n(j)y} \right] = 0. \quad (2)$$

Together, (1) and (2) imply house transaction prices in logs follow a Nonparametric Instrumental Variables (NPIV) model (Newey and Powell, 2003) with linearly separable but nonparametric controls as in Chernozhukov et al. (2018); Robinson (1988). As discussed in Arkhangelsky and Imbens (2023) and Wooldridge (2021), including street $\ell(j)$ and block group $n(j)$ -by-year y fixed effects in a linear model is equivalent to controlling linearly for vectors $\bar{X}_{\ell(j)}$ and $\bar{X}_{n(j)y}$ of street and block group-by-year-level averages of the instrument and observed characteristics, respectively. By instead controlling for group means of higher-dimensional transformations of, and interactions between, monthly street traffic $\tilde{C}_{\ell(j)t(y)}$ and housing unit characteristics A_j and accounting for these controls in a nonparametric fashion, we account for richer forms of unobserved heterogeneity across streets and block groups over time than vanilla fixed effects (Arkhangelsky and Imbens, 2022, 2023)

To estimate the traffic exposure price effect function $p_{t(y)}^{(C)}$ under (1) and (2), we augment an NPIV estimator (see e.g. X. Chen and Christensen (2018)) with a debiased machine learn-

ing approach to controlling for auxiliary variables akin to partially linear IV (Chernozhukov et al., 2018; Robinson, 1988). In particular, we approximate the logarithm of the traffic exposure price effect function $\log p_{t(y)}^{(C)}$ by a linear combination of flexible basis functions and use the elements of a higher-dimensional basis expansion of $\tilde{C}_{\ell(j)t(y)}$ as instruments, in accordance with X. Chen and Christensen (2018). Having reduced effect estimation to a partially linear, multivariate instrumental variables regression, we predict log prices and the terms of the basis expansions of the treatment and instrument in each year using Generalized Random Forests (Athey et al., 2019),¹⁷ subtract those predictions from their realizations as in Chernozhukov et al. (2018); Robinson (1988), and then estimate the coefficients on the treatment basis expansion via two-stage least squares on these residualized variables (Chernozhukov et al., 2018; Robinson, 1988).

3.1.1 Instrumenting for traffic flow changes

To describe our shift-share instrument for traffic flows, it is instructive to first return to how we construct our measures of monthly street car traffic flows. Suppose there are two possible states of the world: “high” traffic H and “low” traffic L that occur with probabilities w_{tL} and $1 - w_{tL}$ in periods $t = 2010$ and $t = 2018$. In each period $t \in \{2010, 2018\}$ and traffic state $s \in \{L, H\}$, every road segment in the road network has a speed, so we can compute the optimal driving routes each household in the Boston CSA would take to get to the POIs they visit in each traffic state. For each period t , road segment ℓ and traffic state s , we let $C_{\ell ts}$ denote the total number of car trips that traverse road segment ℓ each month if all car trips occurred during traffic state s in period t . Given the traffic state probabilities w_{tL} and $1 - w_{tL}$, we can compute the actual number of monthly car trips on road segment ℓ in period t as $C_{\ell t} = w_{tL}C_{\ell tL} + (1 - w_{tL})C_{\ell tH}$.

Recall our concern that changes in traffic flows $C_{\ell t}$ within a neighborhood over time could be correlated with confounding micro-geographic changes in population and/or housing desirability. To address these concerns, we construct a shift-share-like instrument for realized traffic flows in period t with the traffic flows induced *just* by the traffic state probabilities w_{tL} and $1 - w_{tL}$ in that period, while holding the state-specific traffic flows on each road segment fixed in 2010:

$$\tilde{C}_{\ell t} := w_{tL}C_{\ell,2010,L} + (1 - w_{tL})C_{\ell,2010,H}.$$

¹⁷We do not need to cross-fit explicitly since Generalized Random Forest predictions generated via the `grf` package automatically leave out each observation from within-leaf averages when computing in-sample predictions (Athey et al., 2019). Moreover, Random Forests are algorithmically stable (Wang, Wu, and Nettleton, 2023), and thus bias due to over-fitting is minimal when estimating treatment effects in a doubly-robust fashion (Q. Chen, Syrgkanis, and Austern, 2022).

Given our inclusion of road segment-specific “fixed effects” in the estimation strategy described above, this instrument takes advantage of changes in the probability of the low traffic state over time, along with the fact that some streets are exposed to more or less car traffic from this change in state probabilities depending on the extent to which households who drive through a street in the low traffic state change their route in the high traffic state. In this sense, our identifying variation is closest to that discussed in [Borusyak, Hull, and Jaravel \(2023\)](#) and [Borusyak and Hull \(2023\)](#), where our “shares” are the 2010 state-specific traffic flows and our “shifts” are the changes in state probabilities over time. In keeping with the discussion in [Borusyak et al. \(2023\)](#), we assume our controls described in the previous section are rich enough to span the expected instrument.

Since we do not observe these traffic state probabilities, we instead estimate them, as intimated in [Section 2.2](#). In particular, we assume that, for the streets for which we observe traffic speed distributions, the counts of cars TomTom used to construct those speed distributions are proportional to the true traffic counts on those streets.¹⁸ Then, we regress these TomTom sample sizes for road segments ℓ on the vectors of traffic-state-specific counts $C_{\ell s}$ across traffic states to recover these probabilities up to a normalization; see [Appendix B](#) for details. One might be concerned that our estimates of the state probabilities could be contaminated by the same confounding changes in state-specific traffic flows we intend to avoid, so when computing the instrument for a street ℓ , we use state probabilities estimated excluding all road segments in the same Census tract. In practice, we also consider a finer set of 19 different traffic states corresponding to all road segments’ speeds being a particular ventile of their respective speed distributions, not just “low” or “high” traffic; again, see [Appendix B](#) for details.


3.1.2 Results

In [Figure 2](#), we plot the effects in 2018 on housing units at different levels of baseline nearby street traffic of increasing the monthly car trips through a house’s street by the inter-quartile range of the traffic exposure distribution across housing units (plotted in [Figure A.2](#)), or around 27,000 trips. We find that the effect is markedly convex in the affected household’s baseline nearby car trips: moving a house from the 25th percentile of monthly trips down a house’s street to the 75th percentile causes a 20% decrease in home value, while adding the same quantity of trips to the street of a house at the 75th percentile of nearby monthly trips only decreases that house’s value by 2.7%.¹⁹

¹⁸Such an assumption could be violated if the cars from which TomTom derives its traffic speeds are not a representative sample from the population of cars in the Boston CSA.

¹⁹In [Table A.1](#), we also report average effects by quintile of the distribution of housing units’ baseline nearby monthly trips. Since the 95% simultaneous confidence intervals for the average effects on housing

Figure 2: The effect of adding the inter-quartile range of monthly trips on home prices



Notes: We present our NPIV estimates of the 2018 effect on home prices of an increase of own-street traffic equal to the inter-quartile range of monthly trips through streets (around 27,000 trips) by baseline street traffic. Dashed lines are 90% uniform confidence bands computed as in X. Chen and Christensen (2018) via a bootstrap clustered at the street level. We instrument for changes in busyness as described in Section 3.1.1. We plot the distribution of nearby monthly trips across in Figure A.2.

We also consider the effects of an alternative intervention, namely adding 3,000 monthly car trips to households' streets. We view this intervention as an approximation to the effect of adding a 50 unit apartment building to a household's street, where we assume each new resident takes 60 monthly car trips (the median number of monthly trips per driver, per Figure 1). We plot the effects of adding 3,000 trips to the streets of housing units with different baseline levels of nearby monthly car traffic in Figure A.3a, and we report average effects of this intervention on the values of homes in different quintiles of the baseline monthly nearby trip distribution in Table A.1. We find that adding 3,000 monthly trips down the street of a housing unit in the bottom 20% of baseline nearby monthly trips decreases the value of that home by almost 7%, while adding the same quantity of monthly trips to the busiest streets would decrease the values of homes on those streets by an average of 0.035%.²⁰

Our negative house price effects have broadly similar magnitudes to existing negative effects of multifamily housing construction on nearby home prices and rents found in the

units in the first and second quintiles of the nearby monthly trip distribution do not overlap with the 95% simultaneous confidence intervals corresponding to the effects on housing units in the fourth and fifth quintiles, we can reject the null that the average effects are homogenous across baseline monthly trip quintiles at the 5% level.

²⁰We can reject the null hypothesis of homogenous average effects across baseline street traffic quintiles at the 5% level because the 95% simultaneous confidence intervals for these quintile-specific average effects do not overlap.

literature, e.g. [Asquith et al. \(2023\)](#); [Davidoff et al. \(2022\)](#); [Diamond and McQuade \(2019\)](#); [Li \(2022\)](#); [Pennington \(2021\)](#). Although our estimates of adding traffic to the least busy streets are on the higher end of these existing estimates,²¹ these papers mostly study the effects of *realized* multifamily developments on nearby prices and rents, which tend to be sited on busier streets where our estimated effects are the smallest. As such, we view our estimates as being consistent with the negative effects found in this literature.

3.2 Translating price effects into preferences: a hedonic model of residential choices

To translate these treatment effects of increased nearby car traffic on house prices into measures of households’ preferences to avoid nearby car traffic, we interpret our estimates through the lens of a structural, hedonic model of households’ long-run residential choices.²² In particular, we let \mathcal{J} denote the set of housing options available to households in the Boston CSA, and we index individual housing options by j . Each housing option j lies on street $\ell(j)$ in neighborhood $n(j)$ and is endowed with a vector of observable attributes $A_j \in \mathbb{R}^d$, a vertical unobserved quality $\xi_j \in \mathbb{R}$, and a price $P_j > 0$.²³ Each street ℓ exists in a neighborhood $n(\ell)$ and is endowed with two attributes that affect the desirability of living on it: a flow of total monthly cars C_ℓ passing through it and an unobservable variable η_ℓ with unrestricted dimension representing any additional street or neighborhood-specific factors such as views or unobserved amenities.

Household i with wealth W_i chooses a housing option $J_i \in \mathcal{J}$ and aggregate consumption of other goods E_i to maximize an individual-specific, hedonic utility function with the following quasilinear functional form ([Bajari and Benkard, 2005](#); [Berry and Pakes, 2007](#); [Gorman, 1980](#); [Lancaster, 1966](#); [Rosen, 1974](#)):²⁴

$$\begin{aligned} (J_i, E_i) &:= \arg \max_{j \in \mathcal{J}, e > 0} U_{ij}^{(H)} + e \quad \text{such that } e + P_j \leq W_i, \\ U_{ij}^{(H)} &:= -\beta_i C_{\ell(j)} + U_i^{(-C)}(A_j, \eta_{\ell(j)}) + \xi_j. \end{aligned} \tag{3}$$

In (3), β_i denotes household i ’s willingness to pay to reduce the monthly traffic flow $C_{\ell(j)}$

²¹These larger effect estimates are well within the range of effect magnitudes of other neighborhood amenity changes like urban revitalization programs ([Rossi-Hansberg et al., 2010](#)) or the death of street trees ([Han et al., 2024](#)).

²²See [Bishop and Murphy \(2011\)](#) and [Bishop and Murphy \(2019\)](#) for cogent discussions about how to interpret static hedonic residential choice models in the presence of moving costs and uncertainty about future neighborhood characteristics.

²³Since our residential choice model is static, for notational convenience, we drop the time subscripts on variables we used in the previous section.

²⁴This type of model is also referred to as a “pure characteristics” demand model ([Berry and Pakes, 2007](#)).

down their chosen street by one car. As we show in Appendix C.2, our data are inconsistent with essentially any convexity in households’ disutilities from exposure to the traffic on their chosen streets (Agarwal et al., 2023). As such, we restrict our attention to linear disutility from nearby traffic exposure. The function $U_i^{(-C)}$ captures household i ’s utility over the observable and unobservable attributes of housing option j on street $\ell(j)$ in neighborhood $n(j)$. Hedonic models of preferences can represent households’ heterogeneous utilities more realistically than workhorse alternatives like the mixed logit model in the presence of many products (Bajari and Benkard, 2005; Berry and Pakes, 2007). Moreover, the hedonic model also enables convenient estimation of heterogeneous preference parameters that vary flexibly across households, as we will discuss in detail shortly (Bajari and Benkard, 2005).

To map our data on households’ realized housing and POI visit choices into their preference parameters, we first note that, under our model (3), household i gets the following indirect utility from choosing housing option j :

$$U_{ij}^{(H)} = -\beta_i C_{\ell(j)} + U_i^{(-C)}(A_j, \eta_{\ell(j)}) + \xi_j - P_j.$$

Since household i ’s indirect utility from housing option j is determined only by the price P_j of housing option j and the bundle of characteristics

$$X_j := (C_{\ell(j)}, A_j, \eta_{\ell(j)}, \xi_j)$$

with which housing option j is endowed, we can equivalently view household i as choosing a bundle of housing option characteristics X_j from the set of available housing option characteristic bundles \mathcal{X} and paying a price P_j for that bundle (Bajari and Benkard, 2005; Berry and Pakes, 2007; Rosen, 1974). A seminal result by Bajari and Benkard (2005) then dictates that, under no assumptions on the structure of the supply of housing options, if households are maximizing their respective utilities according to (3), there exists a unique, smooth price function $p: \mathcal{X} \rightarrow \mathbb{R}_+$ mapping each housing option’s characteristic bundle X_j to its price P_j , i.e. $P_j = p(X_j)$. Then, we can rewrite household i ’s indirect utility maximization problem as follows:

$$\begin{aligned} \max_{(c,a,\eta,\xi) \in \mathcal{X}} u_i(c, a, \eta, \xi) - p(c, a, \eta, \xi), \\ u_i(c, a, \eta, \xi) := -\beta_i c + U_i^{(-C)}(a, \eta) + \xi, \end{aligned} \tag{4}$$

which is an optimization problem directly over bundles of housing option characteristics.

As illustrated in Figure 5, the set of possible street traffic exposures across over 460 thousand unique streets in the set \mathcal{X} of over 3 million unique housing options is quite rich. As such, we assume households make an essentially continuous choice over nearby street

traffic $C_{\ell(j)}$.²⁵ Then, so long as p is convex, we can identify household i 's willingness to pay to avoid nearby street traffic β_i using the first order condition of their housing choice problem with respect to street traffic (4) (Agarwal et al., 2023; Bajari and Benkard, 2005):

$$\frac{\partial u_i}{\partial c} \Big|_{x=X_{J_i}} - \frac{\partial p}{\partial c} \Big|_{x=X_{J_i}} = 0 \implies \beta_i = \frac{\partial u_i}{\partial c} \Big|_{x=X_{J_i}} = \frac{\partial p}{\partial c} \Big|_{x=X_{J_i}}.$$

In the previous section, we estimated precisely the curvature of the equilibrium house price surface in each period with respect to the monthly car traffic passing by a house $p_t^{(C)}$. As such, we can use derivatives of our estimates of the 2018 equilibrium price surface $p_{2018}^{(C)}$ to back out each household's willingness to pay to avoid each monthly car trip through their home street. We summarize our estimates of these derivatives in Figure A.3b and Table A.1. In Appendix C.1, we describe how we convert our price surface in logs to one in levels by predicting the 2018 values of housing units that did not transact in 2018, as well as how we adjust the price surface slightly to ensure it is consistent with all households' maximizing their respective utility functions.


In Figure 3, we plot the distribution of households' willingnesses to pay to avoid each monthly trip through their home street. We find that the median household is willing to pay \$3.5 in property value to avoid an additional monthly car trip through their chosen home's street, but that this disutility varies by orders of magnitude in the population, with the 10th and 90th percentiles of households willing to pay \$0.05 and \$31.89 per avoided nearby monthly car trip.

4 Valuing households' time costs of congestion

Having characterized households' willingnesses to pay to avoid exposure to nearby street traffic, we now estimate households' heterogeneous preferences over travel time to POIs. We also estimate how those travel times respond endogenously to additional cars on the road. To do so, we build and estimate a structural model of households' POI visiting behavior and road congestion. In our model, households' first allocate monthly time to visiting POIs in different categories and then choose which POIs to visit within each category across space subject to those time constraints. We then integrate this model into our hedonic residential choice framework described in Section 3.2. Intuitively, we are able to infer households' willingnesses to pay for POI access from their realized choices of both POI visits and residential choices.

²⁵Bajari and Benkard (2005) and Agarwal et al. (2023) discuss partial identification of the distribution of preferences across households when \mathcal{X} is considered to be discrete; the richer the choice set \mathcal{X} is, the tighter the bounds on each households' preference parameters are.

Figure 3: The distribution of estimated WTP to avoid trips on own street



Doing so allows us to estimate preferences that are both internally consistent with our hedonic utility framework and have richer variation across households than if we were to calibrate values of time based on existing estimates.

4.1 Model of demand for POI visits

We introduce our model of households’ values for POI visits in reverse order of their decisions. We begin with households’ allocation of time to POI visits given a home neighborhood and then return to how households choose housing options in part based on their differential access to POIs.

4.1.1 Choices of POI visits given residential choices

Having chosen a housing option in a neighborhood, each household chooses how to allocate time within a month to visiting POIs of different types. Let \mathcal{N} denote the set of neighborhoods in the Boston CSA. We assume that in each neighborhood $k \in \mathcal{N}$, there exists a mass M_{ck} of non-atomic POIs in each category c . Each POI ω in neighborhood k and category c provision a differentiated experience with quality $v_{ck}(\omega)$ to the households that choose to visit it. We assume POIs within each neighborhood and category have symmetric fixed and variable costs and compete by choosing their experience quality to maximize profits. In equilibrium, all POIs in the same neighborhood k and category c choose $v_{ck}(\omega)$ equal to the same quality value v_{ck} ; see Appendix D.1 for details.

We model POIs as competing on quality, not prices, in contrast to other models of

POI supply and demand like [Almagro and Dominguez-Iino \(2022\)](#). Our notion of a POI is broader than just consumption amenities like restaurants and bars, encapsulating places like workplaces and schools as well, so not all POI visits have an obvious price associated with them.²⁶ Further, it would be unrealistic to assume POIs compete directly on visit times, since households willingly make frequent, long-duration visits to POIs like workplaces and travel long distances to visit desirable consumption amenities like restaurants and sports venues. As such, allowing POIs within each category to differentiate themselves on quality parsimoniously captures the notion that, in a similar manner across POI categories, there are POIs that households value traveling to despite their high travel times.

Household i with home neighborhood n chooses the quantity of monthly visits to POIs in each category c and neighborhood k to maximize their POI visit utility subject to a time budget constraint. We assume that travel times are stochastic: there exist S different traffic states, each of which corresponds to a distinct profile of car speeds across the roads in the Boston CSA’s road network, and each state s occurs with probability w_s . Under traffic state s , it takes t_{nks} hours to travel from neighborhood n to neighborhood k .²⁷ For computational tractability, we assume that the travel time from every house in neighborhood n to every POI in neighborhood k is constant under traffic state s . Each visit to POI ω in category c and neighborhood k also costs the household the time t_{nks} it takes to travel to and from the POI plus the visit’s dwell time $d_{ck}(\omega)$, which is the time the household spends actually at the POI. Since we have assume POIs are symmetric within a category and neighborhood, we have that dwell times $d_{ck}(\omega)$ are also equal to constants d_{ck} within each category and neighborhood.

Households allocate their time to POI visits in two stages. First, they choose how to allocate their POI visit time to each category of POIs. Second, they choose how much of the time they have allocated to each POI category to spend visiting each POI in that category. We introduce each stage of the household POI visit choice problem in more detail in reverse order.

First, suppose household i with home neighborhood n has already allocated time \tilde{T}_{inc} to visiting category c POIs. We assume they then choose quantities of visits $q_{inck}(\omega)$ to category c POIs in neighborhood k to maximize the following CES utility function subject

²⁶We also do not have access to data on expenditures at POIs either, although such data do exist for POIs at which transactions occur.

²⁷We do not include return travel time in the time cost of a visit, since it is possible the household visits the POI as part of a chain of trips away from their home neighborhood. [Miyachi et al. \(2021\)](#) and [Oh and Seo \(2022\)](#) emphasize the importance of taking trip-chaining into account when modeling POI visits, although according to [Cook \(2022\)](#), they are relatively rare. Our data do not contain the individual-level POI visits necessary to measure chains directly. We leave possible extensions of our model that parsimoniously allow for trip chains to future work.

to total category- c POI visit time being below \tilde{T}_{inc} :

$$U_{inc}^{(\text{POI})}(\tilde{T}_{inc}) := \max_{(\tilde{q}_{inc}(\cdot))_{k \in \mathcal{N}}} \left(\sum_{k \in \mathcal{N}} \int (v_{ck}(\omega) \tilde{q}_{ijck}(\omega))^{\frac{\sigma_c-1}{\sigma_c}} d\mu_{ck}(\omega) \right)^{\frac{\sigma_c}{\sigma_c-1}} \quad (5)$$

such that $\sum_{k \in \mathcal{N}} \int \tilde{q}_{inc}(\omega) \sum_{s=1}^S w_s (t_{nks} + d_{ck}(\omega)) d\mu_{ck}(\omega) \leq \tilde{T}_{inc}$.

Since within each category and neighborhood, POIs are symmetric, the quantities of visits $q_{inck}(\omega)$ that solve the maximization problem in (5) are the same across ω , so we let q_{inck} denote the optimal quantity of visits to each category c , neighborhood k POI chosen by household i living in neighborhood n . By standard CES logic, we have that $\sigma_c > 1$ is the elasticity of time substitution between visits to category c POIs across space.

Next, we describe how household i living in neighborhood n allocates their total POI visit time T_i across POI categories.²⁸ In particular, such a household chooses category time allocations T_{inc} to maximize the following Cobb-Douglas utility function subject to total POI visit time across categories being below T_i , where \mathcal{C} denotes the set of all POI categories:

$$U_{in}^{(\text{POI})} := \max_{(T_{inc})_{c \in [\mathcal{C}]}} \prod_{c \in \mathcal{C}} \left(U_{inc}^{(\text{POI})}(\tilde{T}_{inc}) \right)^{\kappa_{ic}} \quad (6)$$

such that $\sum_{c \in \mathcal{C}} \tilde{T}_{inc} \leq T_i$,

where κ_{ic} are non-negative weights normalized to sum to one that correspond to the shares of total POI visiting time T_i that household i allocates to category- c POI visits. As we will see shortly, homogeneity in inner nest preferences across households enables us to collapse the rich set of POI visiting options into a low-dimensional vector of “sufficient statistics” that represent how accessible POIs of different categories are from a given neighborhood. Heterogeneous outer nest preferences still capture realistic heterogeneity in POI visiting preferences across households.

Under this model, household i in neighborhood n chooses to make the following quantity

²⁸While we do not model how households trade off time spent on POI visits versus other uses of time, one could easily write down an outer Cobb-Douglas utility nest that governs households’ allocation of total time to POI visiting and other uses. Since given T_i , such a decision does not affect household i ’s POI visiting behavior or their choice of residence, we do not explicitly include this decision stage in our model.

of visits to each category c , neighborhood- k POI:

$$q_{inck} = \sum_{s=1}^S w_s (t_{nks} + d_{ck})^{-\sigma_c} v_{ck}^{\sigma_c - 1} I_{cn}^{1-\sigma_c} \kappa_{ic} T_i, \quad (7)$$

$$I_{cn} := \left(\sum_{k' \in \mathcal{N}} M_{ck'} v_{ck'}^{\sigma_c - 1} \sum_{s=1}^S w_s (t_{nk's} + d_{ck'})^{1-\sigma_c} \right)^{\frac{1}{1-\sigma_c}},$$

where I_{cn} is a time cost index measured in hours for visits to category c POIs from home neighborhood n . Intuitively, demand for visiting a POI in a category and neighborhood is increasing in the quality provisioned by POIs in that category and neighborhood and decreasing in the expected time cost of visiting that POI. We have that household i living in neighborhood n 's indirect utility from POI visits is given by

$$U_{in}^{(\text{POI})} = T_i \prod_{c=1}^C \kappa_{ic}^{\kappa_{ic}} I_{cn}^{-\kappa_{ic}}. \quad (8)$$

See Appendix D.2 for derivations of (7) and (8).

4.1.2 Residential choices based on POI access

We now integrate the POI visiting model described above into our hedonic residential choice model (3) introduced in Section 3.2. While in Section C.1, we had no need to restrict household i 's utility $U_i^{(-C)}$ over non-street-traffic housing choice characteristics to estimate their willingness to pay to avoid nearby street traffic, we now decompose $U^{(-C)}$ into two additive components. The first is a neighborhood-specific POI visit utility component, and the second is a component corresponding to observable housing option characteristics and street-and-neighborhood unobservables. Concretely, we write household i 's utility from living in housing option j can be written as follows:

$$U_{ij}^{(\text{H})} := -\beta_i C_{\ell(j)} + \gamma_i \log(U_{in}^{(\text{POI})}) + U_i^{(\text{A})}(A_j, \tilde{\eta}_{\ell(j)}) + \xi_j.$$

Substituting the expression for household i 's indirect POI visiting utility into the expression above, we have that household i gets the following indirect utility from choosing housing option j :

$$U_{ij}^{(\text{H})} = K_i - \beta_i C_{\ell(j)} - \gamma_i \sum_{c \in \mathcal{C}} \kappa_{ic} \log(I_{n(j)c}) + U_i^{(\text{A})}(A_j, \eta_{\ell(j)}) + \xi_j - P_j, \quad (9)$$

where K_i is a collection of terms that don't depend on which housing option household i chooses.²⁹

In (9), household i 's indirect POI visiting utility from housing option j is determined completely by the bundle of travel time indices $\{I_{n(j)c}\}_{c \in \mathcal{C}}$ that characterize POI access from housing option j 's neighborhood $n(j)$. As such, we can view households as making direct choices over a low-dimensional ‘‘sufficient statistic’’ of a neighborhood’s POI access, not intractably high-dimensional visit counts to every POI in the Boston CSA. This choice problem simplification allows us to estimate households’ willingness to pay for POI access hedonically. We rewrite household i 's hedonic indirect utility maximization problem directly over housing choice characteristics (4) as follows:

$$\begin{aligned} & \max_{(c, \iota, a, \eta, \xi) \in \mathcal{X}} u_i(c, \iota, a, \eta, \xi) - p(c, \iota, a, \eta, \xi), \\ u_i(c, \iota, a, \eta, \xi) & := K_i - \beta_i c - \gamma_i \sum_{c \in \mathcal{C}} \kappa_{ic} \log(\iota_c) + U_i^{(A)}(a, \eta, \xi). \end{aligned}$$

Since there are 5,887 Census block groups in the Boston CSA, we assume households’ choices over travel time index bundles are also effectively continuous. Then, given travel time indices I_{nc} for every category and neighborhood recovered using our POI visiting data, we can identify household i 's taste for POI visits γ_i from household i 's first-order conditions with respect to $\log(\iota_c)$:

$$\begin{aligned} & \sum_{c \in \mathcal{C}} \left[\frac{\partial u_i}{\partial \log(\iota_c)} \Big|_{x=\mathcal{X}_{J_i}} - \frac{\partial p}{\partial \log(\iota_c)} \Big|_{x=\mathcal{X}_{J_i}} \right] = 0 \\ \implies \gamma_i & = \sum_{c \in \mathcal{C}} \frac{\partial u_i}{\partial \log(\iota_c)} \Big|_{x=\mathcal{X}_{J_i}} = \sum_{c \in \mathcal{C}} \frac{\partial p}{\partial \log(\iota_c)} \Big|_{x=\mathcal{X}_{J_i}}. \end{aligned} \tag{10}$$

Given this result, we now turn to estimating our travel time indices along with the curvature of the equilibrium price surface p with respect to these travel time indices.

4.2 Travel time preference estimation

Estimation of households’ preferences for visiting POIs proceeds in two steps. First, we estimate the parameters of the POI visiting model that determine the travel time indices I_{nc} across each of the seven POI categories introduced in Section 2, as well as households’ heterogeneous preferences for different POI categories κ_{ic} . Second, we estimate the derivatives of the equilibrium price surface with respect to each travel time index and aggregate them into estimates of γ_i according to (10).

²⁹Specifically, $K_i := \log(T_i) + \sum_{c \in \mathcal{C}} \kappa_{ic} \log(\kappa_{ic})$.

4.2.1 POI visiting model components

To estimate the parameters of the POI visiting model described in Section 4.1, we first note that, in our SafeGraph data, we observe noisy estimates \hat{Q}_{nck} of the 2018 total quantity of visits $M_{ck} \mathbb{E}[q_{in(J_i)ck} \mid n(J_i) = n]$ by the average resident of neighborhood n to all POIs in category c and neighborhood k , as well as the number M_{ck} of category c POIs in neighborhood k and the travel t_{nks} and dwell d_{ck} times required to visit that POI in each traffic state s . Assuming \hat{Q}_{nck} is an unbiased estimate of Q_{nck} , we can use (7) to derive a conditional moment restriction that identifies our travel time elasticities σ_c :

$$\begin{aligned} \mathbb{E}[\hat{Q}_{nck} \mid M, v, t, d] &= M_{ck} \mathbb{E}[q_{in(J_i)ck} \mid n(J_i) = n, v, t, d] \\ &= \underbrace{M_{ck} v_{ck}^{\sigma_c - 1}}_{\exp(\nu_{ck})} \underbrace{I_{nc}^{1 - \sigma_c} \mathbb{E}[\kappa_{ic} T_i \mid n(J_i) = n]}_{\exp(\delta_{cn})} \sum_{s=1}^S w_s (t_{nks} + d_{ck})^{-\sigma_c} \\ &= \exp(\nu_{ck} + \delta_{cn}) \sum_{s=1}^S w_s \exp(-\sigma_c \log(t_{nks} + d_{ck})). \end{aligned} \quad (11)$$

(11) resembles the conditional moment restriction corresponding to a category-specific Poisson regression with two-way fixed effects, albeit with a mixture over the terms governing responsiveness to travel times under different traffic states. We implement a computationally efficient algorithm to estimate the within-category travel time elasticities σ_c and the high-dimensional fixed effects ν_{ck} and δ_{cn} that resembles the highly performant fixed-effects Poisson regression algorithm underlying the `fixest` R package (Bergé, 2018). Given these parameter estimates, we can back out the visit qualities v_{ck} up to a normalization from our estimates of the travel time elasticities σ_c and fixed effects ν_{ck} . Recovering these parameters allows us to construct estimates of the travel time indices I_{nc} . We provide details in Appendix D.3.

To recover households' heterogeneous preferences over POI categories, we assume all households that choose to live in a Census block group have the same preferences for time spent visiting POIs in each category. While such an assumption masks within-neighborhood heterogeneity in households' preferences, it still allows for rich variability in households preferences across the 5,887 Census block groups in the Boston CSA. Under this assumption, we can estimate the time $\hat{T}_{in(J_i)c}$ household i choosing to live in neighborhood $n(J_i)$ allocates to visiting category- c POIs via a function of our existing parameter estimates; dividing each by their sum across categories yields estimates of κ_{ic} for each category. We provide more details in Appendix D.3. In Figure A.4, we visualize the spatial distributions of κ_{ic} for each category across Census block groups.

4.2.2 Hedonic willingness to pay for POI access

Armed with estimates of the travel time indices I_{nc} characterizing accessibility of POIs in each category c from each neighborhood n , we proceed to identify and estimate households' willingnesses to pay for POI access. To do so, we use a similar identification strategy as we did in Section 3.1 to identify the causal effects of streets' monthly traffic on the prices of homes on those streets.

In particular, we use our POI visit model parameter estimates to construct period-specific travel time indices I_{nct} , where we allow traffic state-specific travel times and the probabilities of those states to correspond to the period t .³⁰ We then estimate how changes, across Census block groups, in category-specific POI access between 2010 and 2018 affect home transaction prices after removing the effect of changes in own-street traffic exposure.³¹ Again, to address the concern that changes in POI access within a neighborhood might be correlated with unobserved changes in neighborhood desirability, we instrument for the travel time indices I_{nct} similarly to how we did in Section 3.1.1, namely with versions of the travel time indices \tilde{I}_{nct} that hold traffic state-specific travel times fixed at their 2010 levels and only allow traffic state probabilities to change over time.³²

In practice, we impose additional functional form restrictions on the equilibrium price surface for tractability. We decompose the non-traffic-exposure term $p^{(-C)}$ in (1) into a product of unrestricted functions of each POI category-specific travel time index, as well as an unrestricted function $p^{(A)}$ of housing unit j 's attributes, as well as street $\ell(j)$ and municipality $m(j)$ -by-year y unobservables:

$$P_{jy} = \exp(\xi_{jy}) p_{t(y)}^{(C)}(C_{\ell(j)t(y)}) p_y^{(A)}(A_j, \phi_{\ell(j)}, \varphi_{m(j)y}) \prod_{c \in \mathcal{C}} p_{ct(y)}^{(I)}(I_{n(j)ct(y)}).$$

Now, ξ_{jy} encapsulates unobserved, time-varying factors that affect housing option desirability that differ at levels up to Census block groups.

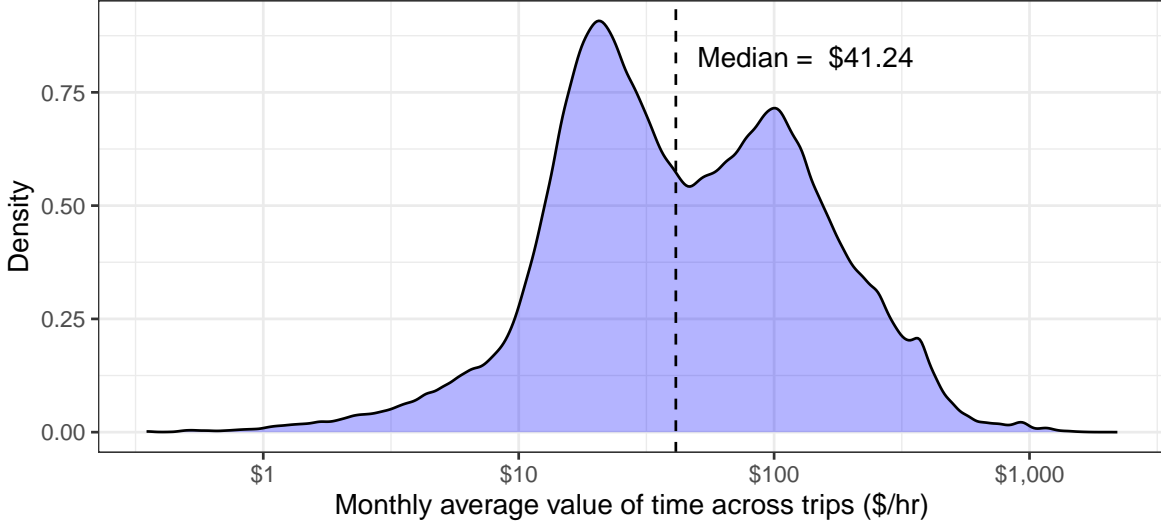
To estimate the functions $p_{ct(y)}^{(I)}$, we first divide each observed transaction price by our

³⁰Because we only observe a single cross-section of POIs in 2018, when constructing these travel time indices in each year, we hold the numbers of POIs in each neighborhood and category their visit qualities, and their dwell times fixed at their 2018 values.

³¹We cannot directly relate prices with POI-visiting indirect utility $U_{in(j)}^{(\text{POI})}$ directly because POI-visiting utility is comprised of both households' homogenous preferences for POI access within each category *and* their heterogeneous preferences across categories. However, [Bajari and Benkard \(2005\)](#)'s equilibrium price surface existence theorem requires the price surface to depend on product characteristics that are experienced in the same way by all households. Because households choose to live in locations that give them more POI visiting utility given their preferences, the relationship between house prices and indirect POI visiting utility will be dampened.

³²Again, we use probabilities estimated using only roads outside the Census tract containing block group $n(j)$.

Figure 4: The estimated distribution of average value of travel time per POI visit



Notes: We present the distribution across households of their values of an additional hour of travel time on average over their visits to POIs, as computed based on the heuristic procedure described in Appendix D.4.

estimate of the street traffic exposure term $p_{t(y)}^{(C)}(C_{\ell(j)t(y)})$ from Section 3.1. Then, we implement a similar nonparametric instrumental variables estimator with partially linear controls to the estimator described in Section 3.1. We approximate the logarithm of each category-specific price function term $\log p_{ct(y)}^{(I)}$ with a linear combination of flexible basis functions of $I_{n(j)ct(y)}$, and we instrument for these variables with higher-dimensional basis expansions of our instruments $\tilde{I}_{n(j)ct(y)}$ (X. Chen and Christensen, 2018). We then residualize prices and these travel time index and instrument basis expansions using Generalized Random Forest-predicted versions of these variables based on housing unit attributes as well as street and municipality-by-year-level averages of the instruments and housing unit characteristics for transacted units (Athey et al., 2019; Chernozhukov et al., 2018). As discussed in detail in Section 3.1, doing so controls for time-invariant, street specific and time-varying, municipality-specific unobserved heterogeneity (Arkhangelsky and Imbens, 2022, 2023). We estimate the coefficients on the travel time index basis expansions via two-stage least squares using the residualized prices, treatments, and instruments (X. Chen and Christensen, 2018; Chernozhukov et al., 2018). Finally, we translate these price surface component estimates into estimates of hedonic POI-visiting preferences γ_i by adding up derivatives of our estimates of $p_{ct(y)}^{(I)}$ across categories.

To summarize our heterogeneous preference estimates in an interpretable way, we compute rough estimates of households’ values’ of an hour of travel time on average across their visited POIs under our model, which we plot in Figure 4; see Appendix D.4 for details. We

find that the median household’s value of travel time to POIs is around \$41 per hour, with the 10th and 90th percentiles of the distribution across households being \$12 and \$199 per hour. To benchmark these estimates, we divide each household’s value of time estimate by the median hourly household wage in that household’s chosen Census block group and plot the resultant distribution in Figure A.5. We find that the median household has a value of time at around 110% of their neighborhood hourly median wage, with 10th, 25th, 75th, and 90th percentiles of 27%, 48%, 271%, and 511%, respectively. As discussed in Section 1,³³ these estimates are broadly in line with the range of estimates from the transportation and spatial economics literatures. However, our increased preference heterogeneity relative to many existing approaches means that we find some households’ values of time that fall slightly outside of the range of existing estimates.

4.3 Travel time effects of increased traffic volume

To translate increases in car trips on roads into welfare costs of congestion, we estimate the empirical relationship between cars’ speeds on roads and the volume of traffic on those roads. To do so, we follow Yang et al. (2020) and a rich literature dating back to Greenshields, Bibbins, Channing, and Miller (1935) and posit a linear structural relationship between traffic speed on a road and traffic *density* on that road, defined as the number of cars per lane per kilometer and equivalent to traffic volume divided by traffic speed. We do so since, in equilibrium, the same traffic volume can correspond to high or low traffic speed (Hall, 1996).³⁴

In particular, we define monthly average car speed $V_{\ell t}$ and monthly average car density $D_{\ell t}$ for road segment ℓ in period t as the traffic-state-probability-weighted averages of state-specific monthly traffic speeds $V_{s\ell t}$ and densities for road segment ℓ in year t :³⁵

$$V_{\ell t} := \sum_{s=1}^S w_{st} V_{s\ell t}, \quad D_{\ell t} := \sum_{s=1}^S w_{st} \frac{C_{s\ell t}}{\lambda_{\ell} V_{s\ell t}} \cdot \frac{1.36 \times 10^{-3} \text{ months}}{\text{hour}}, \quad (12)$$

where λ_{ℓ} denotes road segment ℓ ’s number of lanes and we assume that monthly traffic volume in a given state $C_{s\ell t}$ is spread uniformly across lanes and time spent in that state per month. Then we assume that average traffic speed $V_{\ell t}$ and density $D_{\ell t}$ are related via

³³See Section 5.2.2 of Su (2022a) for a thorough review of different literatures’ estimates of similar quantities. Because these literatures use different approaches for estimating value-of-time-like quantities, unified interpretation across approaches is difficult.

³⁴See Section I in Yang et al. (2020) and Appendix D.2 in Cook and Li (2023) for in-depth explanations of this fact and the value of estimating congestion effects in terms of traffic density, not traffic volume directly.

³⁵There are 730 hours per month, which means there are $1/730 = 1.36 \times 10^{-3}$ months per hour.

the following linear function:

$$V_{\ell t} = \underbrace{(b_0 + b_v \bar{V}_\ell)}_{b_\ell} D_{\ell t} + \underbrace{(a_\ell + a_{r(\ell)t} + \zeta_{\ell t})}_{a_{\ell t}}, \quad (13)$$

where we allow the slope b_ℓ to vary with ℓ 's speed limit \bar{V}_ℓ and assume the free-flow speed $a_{\ell t}$ (speed in the absence of any cars) is determined by the sum of a road segment fixed effect, road type by period fixed effect, and a residual.

Estimating the regression specification (13) assuming $\zeta_{\ell t}$ are uncorrelated with $D_{\ell t}$ given fixed effects ignores the possibility that changes in speed on a road segment over time could be correlated with changes in the desirability of driving on that road in equilibrium (e.g. decreasing speed on a road leading fewer cars to use that road) (Akbar et al., 2023; Couture et al., 2018; Yang et al., 2020). To overcome this challenge, we instrument for changes in $D_{\ell t}$ over time using just the changes over time in traffic density due just to changes in the traffic state probabilities and holding traffic density fixed in 2010, as in our hedonic valuations of car traffic exposure and travel time earlier.

We then translate these effects of traffic density on speeds into effects of traffic volumes on speeds similarly to Yang et al. (2020). Because velocity appears in the definition of density (12), we solve each traffic states-specific congestion function for endogenous speed $V_{\ell s}$, choosing the solution in which traffic speed is decreasing in traffic volume:³⁶

$$V_{\ell s}(c) = a_\ell + (1.36 \times 10^{-3})b_\ell \frac{c}{\lambda_\ell V_{\ell s}(c)} \implies V_{\ell s}(c) = \frac{1}{2} \left(a_\ell + \sqrt{a_\ell^2 + (5.44 \times 10^{-3})b_\ell \frac{c}{\lambda_\ell}} \right).$$

The inverse of $V_{\ell s}$ times the length L_ℓ of road segment ℓ yields the time it takes to traverse road segment ℓ .


In Figure 5, we plot the distribution of 2018 traversal time elasticities with respect to monthly traffic volume across the Boston CSA's road segments. These elasticities are small in general, although they are larger than those implied by the congestion function estimates from Yang et al. (2020).

5 Counterfactuals

Having estimated households' preferences and roads' congestion technology, we now study the heterogeneous traffic costs incurred by incumbent residents from adding additional households to different locations across the Boston CSA. We then use these location-specific welfare

³⁶As explained in Hall (1996), Yang et al. (2020), and Cook and Li (2023), traffic speed can be increasing in traffic volume if traffic is so backed up that increasing speeds increases throughput of cars.

Figure 5: The estimated distribution of congestion effects



impacts to evaluate the traffic externalities caused by targeting a 10% increase in metropolitan area population according to several heuristics akin to proposals by policy makers in other metropolitan areas across the United States.

5.1 Simulating a new household’s traffic externalities

Adding a new resident to a street who drives does not just affect the other households living on that street. Anyone living on the streets on which that new resident drives will be exposed to that new resident’s car trips, and any households driving on the same streets will also be slowed down by congestion caused by the new resident’s car trips. To simulate these traffic externalities caused by adding a counterfactual new housing unit to a street ℓ , we make several simplifying assumptions. First, we assume that the new household moving into the new unit takes the same car trips to POIs as the other households living on that street. Given that a household selecting to move into this counterfactual housing unit would likely do so in part based on the POIs that are most convenient to visit from ℓ , such an assumption is not unreasonable. However, if the housing unit were materially different from the existing housing stock on street ℓ , e.g. a new apartment on a street with only single-family homes, it might attract a household with different preferences for visiting POIs.

Second, we assume that the new household travels via public transit at the same rate as incumbent residents on street ℓ . Recall from Section 2 that we use the share of adults in ℓ ’s Census block group that commute by car as our measure of the probability a given household creates car traffic. As above, such an assumption is reasonable so long as the arriving household’s taste for the new housing unit’s characteristics is not strongly correlated

with their propensity to travel via public transit.

Third, we take a first-order approach to welfare analysis: we assume that in response to a marginal increase in car trips through roads, incumbent households do not move, POI quantities and visit qualities do not adjust as a result of the counterfactual housing unit being added, and households do not adjust their optimal, traffic state-specific car routes from their homes to the POIs they visit. Given that the median household visits 60 POIs per month and the median household lives on a street with around 7,000 monthly car trips, such assumptions appear to be reasonable. However, when we study variants of Massachusetts’s MBTA Communities law that require municipalities to plan for a 10% increase in metropolitan area population in Section 5.3, we will view our estimates of the traffic externalities created by these policies as approximate extrapolations.

To define our measures of the welfare effects of adding a new household to a street ℓ , let $\Delta C_{\ell\ell's}$ denote the number of monthly car trips taken by a new household on street ℓ that pass through street ℓ' under traffic state s as they visit POIs. In addition, let N denote the total number of households in the Boston CSA, and for a given partition of those households into disjoint demographic groups, e.g. households in different income quartiles, let G_i denote the demographic group to which household i belongs. We define the total welfare cost to households in demographic group g from exposure to the additional car trips on their streets generated by a new household on street ℓ as follows:

$$\Delta W_{\ell g}^{(E)} := N \mathbb{E} \left[\mathbb{1}\{G_i = g\} \cdot \beta_i \sum_{s=1}^S w_s \Delta C_{\ell\ell(J_i)} \right].$$

The welfare cost consists of adding up households in group g ’s welfare costs per additional monthly trip multiplied by the number of monthly trips taken by the new household on street ℓ that pass through affected households’ chosen streets under different traffic states.

To define the congestion welfare cost from a new household on street ℓ ’s car trips due to increases in travel times, let \mathcal{R}_{nks} denote the sequence of road segments comprising the fastest route traversed by a household traveling from Census block group n to Census block group k under traffic state s . Then, we can define the total first-order travel time welfare cost on households in group g caused by a new household on street ℓ as follows:

$$\Delta W_{\ell g}^{(T)} := N \mathbb{E} \left[\mathbb{1}\{G_i = g\} \cdot \underbrace{\gamma_i \sum_{c \in \mathcal{C}} \kappa_{ic} \sum_{s=1}^S \frac{\partial \log(I_{nc})}{\partial t_{nks}}}_{\uparrow \text{travel time effect on } i\text{'s utility}} \underbrace{\sum_{\ell' \in \mathcal{R}_{nks}} \frac{-L_{\ell'}}{V_{\ell's}^2} \frac{\partial V_{\ell's}}{\partial c}}_{\uparrow \text{trips' effect on travel time from } n \text{ to } k} \bigg|_{c=C_{\ell'}} \Delta C_{\ell\ell'} \right].$$

Here, we combine our estimates of how each road segment’s traversal time is affected by an increase in car trips down that road segment with our estimates of households’ preferences to

minimize time spent visiting POIs. Finally, we define the total traffic externalities incurred by incumbents in demographic group g due to adding a new household on street ℓ as the sum of the two welfare costs defined above:

$$\Delta W_{\ell g} := \underbrace{\Delta W_{\ell g}^{(E)}}_{\text{total exposure welfare costs}} + \underbrace{\Delta W_{\ell g}^{(T)}}_{\text{total congestion welfare costs}} .$$

5.2 Traffic externalities from new households across space

We report statistics characterizing the distributions of traffic exposure and congestion externalities across potential locations for new households for both all households and different household income quartiles in Table 1 and plot these distributions in Figure A.6. We find that the median new resident causes \$11,127 in traffic externalities. However, that quantity varies substantially across space: the 10th and 90th percentiles of the distribution of total traffic externalities caused by a new resident are \$5,553 and \$21,795, respectively. Exposure externalities also tend to vary more than congestion externalities across potential housing locations: the standard deviation of exposure externalities across potential new housing locations is 7,569, while the standard deviation of congestion externalities is just \$2,563. The relative sizes of exposure and traffic externalities also do not vary in lock step across space. In particular, 25% of potential additional residents would generate exposure externalities at most half the size of the congestion externalities they would generate, while almost 25% would generate exposure externalities at least twice the size of the congestion externalities they would generate. These patterns are also broadly replicated when considering the welfare impacts on households living in different quartiles of the household median block group income distribution.

5.3 Evaluating metropolitan-area-wide planning policies’ traffic costs

Next, we aggregate these per-household welfare effects across locations into estimates of the traffic externalities caused by variants of Massachusetts’ MBTA Communities law. Passed in 2021, the law mandated that the 177 municipalities besides Boston served by the Massachusetts Bay Transit Authority (MBTA) modify their land use regulations to enable a cumulative 10% increase in the Boston CSA’s housing stock.³⁷ In particular, each municipality served by the MBTA, or adjacent to an MBTA community, was allocated a number of additional housing units they had to accommodate within their borders via changes to their zoning code. These allocations were determined based on how central a municipality is in the

³⁷Boston was supposedly exempted from the law since it had already made efforts to adjust its zoning code to enable the construction of more housing (Mass Live, 2024).

Table 1: Heterogeneity in a new household's traffic externalities across locations

Welfare cost type	Mean	Std. dev	Quantiles				
			10	25	50	75	90
All households							
Total	\$13,007	\$8,342	\$5,553	\$7,917	\$11,127	\$15,579	\$21,795
Exposure	\$7,359	\$7,569	\$1,790	\$2,884	\$5,112	\$8,976	\$14,838
Congestion	\$5,636	\$2,563	\$2,359	\$3,728	\$5,557	\$7,344	\$8,972
Exposure / Congestion	165%	229%	34%	55%	100%	187%	336%
Income: Bottom 25%							
Total	\$1,252	\$1,333	\$510	\$677	\$961	\$1,333	\$1,978
Exposure	\$637	\$1,288	\$108	\$177	\$309	\$548	\$1,177
Congestion	\$612	\$291	\$282	\$407	\$567	\$762	\$996
Exposure / Congestion	129%	290%	18%	30%	54%	109%	249%
Income: 25% to 50%							
Total	\$2,635	\$3,518	\$873	\$1,169	\$1,585	\$2,427	\$5,073
Exposure	\$1,656	\$3,473	\$189	\$289	\$498	\$1,231	\$4,089
Congestion	\$973	\$459	\$445	\$672	\$915	\$1,207	\$1,563
Exposure / Congestion	245%	627%	19%	30%	55%	158%	567%
Income: 50% to 75%							
Total	\$3,732	\$3,747	\$1,105	\$1,745	\$2,589	\$4,066	\$7,611
Exposure	\$2,180	\$3,587	\$207	\$386	\$773	\$2,250	\$5,854
Congestion	\$1,544	\$718	\$620	\$1,022	\$1,502	\$1,992	\$2,474
Exposure / Congestion	171%	346%	16%	27%	56%	163%	432%
Income: Top 25%							
Total	\$5,103	\$5,402	\$871	\$1,546	\$3,164	\$6,581	\$11,996
Exposure	\$2,625	\$4,516	\$64	\$170	\$618	\$2,948	\$8,033
Congestion	\$2,470	\$1,516	\$722	\$1,233	\$2,214	\$3,510	\$4,603
Exposure / Congestion	93%	179%	7%	12%	28%	99%	255%

Notes: This table reports summary statistics of the distributions of welfare costs on incumbent residents from the traffic externalities caused by adding a new household to different locations across the Boston CSA. Each potential new household location is weighted by the total amount of residential land along focal street on which the new housing unit would be sited. For density plots of each type of welfare cost, see Figure A.6.

MBTA commuter rail network and the size of their existing housing stock ([Massachusetts Executive Office of Housing and Livable Communities, n.d.](#)). Importantly, these zoning code changes were required to be made solely on land within half a mile of a public transit stop.³⁸ We present these allocations of units to municipalities along with eligible land within half a mile of a public transit stop in Figure A.7.

We evaluate how consequential the decision by the Massachusetts State Legislature to concentrate this increase in planned density near public transit was for incumbent residents of the Boston CSA who have otherwise fought to avoid meaningful increases in density ([Boston Globe, 2021](#)). To do so, we use our estimates of the traffic externalities caused by new households from Section 5.2 to compute the total traffic externalities incurred by incumbents if the housing units the MBTA Communities law aimed to encourage were actually built. We then compare those costs to the total traffic externalities caused by a counterfactual variant of the law that would instead spread each municipality’s allocated units across all residential land in that municipality. Finally, we also evaluate another counterfactual policy variant where we allocate each municipality’s housing units exclusively to residential land along non-residential streets.³⁹

In particular, for each policy variant, we assume that each municipality’s allotted units are spread across residential land within their borders and within the area targeted by the policy variant to achieve the minimum increase in density necessary to accommodate those new units. In Figure A.8, we map these minimum density increases by Census block group across the Boston CSA. Having allocated units to locations across the Boston CSA, we add up the welfare effects computed in Section 5.2 created by each of these new units and report our total traffic externality tabulations in Table 6.

We find that the MBTA Communities law’s allocation of housing units to locations near public transit stops would cause \$3.3 billion in traffic externalities. The law’s welfare impacts are \$820 million smaller than the \$4.1 billion in welfare costs from traffic caused by the policy that would densify all residential land. \$500 million (64%) of these savings are driven by lower exposure externalities. For both policies, around 65% of the total welfare cost is caused by exposure externalities.

While this targeting of increased density to locations near public transit reduces total traffic externalities from both exposure and congestion, it does so in a regressive fashion. Compared to the uniform density increase policy variant, building near public transit in-

³⁸The law also required that at least one of the zoning districts created to satisfy the law allow multi-family housing to be built by right ([Massachusetts Executive Office of Housing and Livable Communities, n.d.](#)).

³⁹In the parlance of OpenStreetMap, we allocate housing units to residential land along roads with the “Secondary” classification, e.g. Broadway Street in Cambridge, Massachusetts, or larger.

Figure 6: Traffic costs of MBTA Communities law variants

Density increase locations	Welfare cost (millions)		
	Total	Exposure	Congestion
All households			
Everywhere	\$4,094	\$2,633	\$1,458
Near transit	\$3,273	\$2,133	\$1,137
On thoroughfares	\$2,696	\$1,287	\$1,408
Income: Bottom 25%			
Everywhere	\$452	\$278	\$173
Near transit	\$551	\$403	\$148
On thoroughfares	\$386	\$217	\$169
Income: 25% to 50%			
Everywhere	\$655	\$429	\$225
Near transit	\$547	\$365	\$182
On thoroughfares	\$462	\$243	\$220
Income: 50% to 75%			
Everywhere	\$1,011	\$618	\$392
Near transit	\$758	\$456	\$301
On thoroughfares	\$661	\$282	\$378
Income: Top 25%			
Everywhere	\$2,003	\$1,212	\$788
Near transit	\$1,462	\$845	\$614
On thoroughfares	\$1,298	\$537	\$756

creases the traffic externalities borne by households living in the lowest-income quartile of Census block groups by 22%, while it decreases the traffic externalities borne by households living in the highest-income quartile of Census block groups by 27%. This increased cost to lower income households is driven entirely by a \$125 million increase in exposure externalities.

More optimistically, we find that targeting increased density to residential land along thoroughfares only causes \$2.7 billion in traffic externalities, an additional \$577 million decrease in welfare costs over the original MBTA Communities law and a cumulative savings of \$1.4 billion in traffic externalities relative to increasing density across each municipality. These savings are almost entirely driven by substantial decreases in exposure externalities, as adding new units on thoroughfares causes only marginally smaller congestion externalities than adding those units everywhere and therefore somewhat larger congestion externalities than adding those units near public transit stops. Crucially, this density targeting policy

decreases total traffic externalities for households living in Census block groups at all parts of the income distribution relative to both alternative policy variants.

Finally, we consider a version of the MBTA Communities law’s allocations of new units to municipalities that also requires Boston to increase its density. We allocate units to Boston in accordance with the law’s guidelines and scale down other municipalities’ obligations proportionally to ensure the same total number of units get built as in our other counterfactuals. In Table A.9, we report the total traffic externalities caused by the different targeting rules we considered before under this alternative allocation of units to municipalities that includes units in Boston. We find that shifting new housing units towards Boston increases total welfare costs slightly from traffic under all three targeting policies. Perhaps surprisingly, despite fewer new trips congesting roads leading into Boston, total congestion externalities only decrease by tens of millions of dollars across all three targeting policy variants, and these savings are offset by increases in exposure externalities. Another important difference is that, because much of the residential land in Boston is within half a mile of a public transit stop (as can be seen in Figure A.7), the welfare gains from building near transit relative to uniformly are muted by the fact that the targeting policies’ welfare costs in Boston are similar. As such, the gap in welfare costs between building new units near transit as stipulated by the original MBTA communities law and building those new units on thoroughfares is larger than it was when using the original MBTA Communities law’s unit allocations. Otherwise, the broad incidence patterns across income groups are similar to those under the counterfactuals based on the original law’s unit allocations.

Before concluding, we discuss several important caveats to this analysis. First, as discussed in Section 5.1, we do not account for households changing their chosen driving routes in equilibrium, which would dampen the travel time costs for exposed households but expose additional households to congestion on chosen alternative routes. We also do not accommodate nonlocal movement along each street’s congestion function in response to increases in traffic volume, which, at least under the concave congestion function functional form assumed in (13), would lead to larger welfare impacts. More broadly, we do not account for general equilibrium responses to increases in population density like shifts in the locations of and visit qualities provisioned by POIs across the Boston CSA. Increased demand for POIs could spur POI entry near new residents, concentrating their car traffic more locally and changing welfare effects in ways that are difficult to predict. While accounting for these additional forces would yield more realistic predictions of the effects of these policies, we leave doing so for future work.

6 Conclusion

In this paper, we measure the spatially heterogeneous costs of traffic exposure and congestion to households. We construct estimates of monthly car trip counts through every street in the Boston metropolitan area using GPS data on household POI visits and households’ optimal driving routes. We then take a hedonic approach to valuing households’ exposure to car traffic on their streets, estimating the causal effect of increases in nearby traffic volume on home transaction prices. Our nonparametric estimates indicate substantial heterogeneity in households’ disutilities from traffic exposure, with the median household willing to pay around \$3.50 to avoid an additional monthly car trip on their street.

We also estimate households’ heterogeneous distastes for travel times to different categories of POIs using a structural model of households’ POI visit decisions. Under this model, a household’s indirect utility from a neighborhood depends on a household-specific weighted sum of neighborhood-level, category-specific “travel time indices” that capture the visit-quality-adjusted difficulties of accessing POIs of each type. We again use nonparametric instrumental variables estimates to recover the causal effects of changes in these travel time indices on house prices; translating these effects into households’ preferences shows that households’ values of time also vary substantially.

Given estimates of households’ preferences and the effects of traffic volume on road speeds, we simulate the heterogeneous welfare costs to incumbent residents from adding a new resident and their associated car traffic to each location across the Boston metropolitan area. We find that the traffic externalities caused by adding one new resident vary substantially: the gap between the 10th and 90th percentiles is over \$15,000. Exposure costs account for a majority of these externalities in many locations. Finally, we use our estimates to evaluate the traffic externalities caused by variants of Massachusetts’ 2021 MBTA Communities law, which required municipalities to increase housing density near public transit stations. We find that targeting upzoning near public transit as specified in the law reduces externalities by around \$820 million compared to uniform density increases, but that targeting development to thoroughfares instead could save an additional \$580 million.

Throughout this paper, we focused on the traffic externalities caused by adding new residents, we could use this same framework in future work to evaluate other planning decisions made by policy makers. For example, we could combine our per-household welfare estimates with a richer transit mode choice model to evaluate the traffic externality reductions induced by new public transit infrastructure across space, or we could evaluate the traffic externality reductions from locating POIs closer to the households that visit them most. Both of these interventions have the potential to mitigate the traffic externalities caused

by new housing construction, and studying them could highlight complementarities between planning policies that encourage new residents and policies intended to minimize their costs to incumbents.

References

- Agarwal, N., Li, P. Z., and Somaini, P. J. (2023). *Identification using revealed preferences in linearly separable models* (Tech. Rep.). National Bureau of Economic Research. [8, 19, 20, 56, 57]
- Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6), 2127–2189. [8]
- Akbar, P. A., Couture, V., Duranton, G., and Storeygard, A. (2023). *The fast, the slow, and the congested: Urban transportation in rich and poor countries* (Tech. Rep.). National Bureau of Economic Research. [6, 30]
- Allen, T., and Arkolakis, C. (2022). The welfare effects of transportation infrastructure improvements. *The Review of Economic Studies*, 89(6), 2911–2957. [7]
- Almagro, M., Barbieri, F., Castillo, J. C., Hickok, N. G., and Salz, T. (2024). *Optimal urban transportation policy: Evidence from chicago* (Tech. Rep.). National Bureau of Economic Research. [7]
- Almagro, M., and Dominguez-Iino, T. (2022). Location sorting and endogenous amenities: Evidence from amsterdam. Available at SSRN 4279562. [8, 9, 22]
- Anagol, S., Ferreira, F. V., and Rexer, J. M. (2021). *Estimating the economic value of zoning reform* (Tech. Rep.). National Bureau of Economic Research. [8]
- Arkhangelsky, D., and Imbens, G. W. (2022). Doubly robust identification for causal panel data models. *The Econometrics Journal*, 25(3), 649–674. [4, 13, 14, 28]
- Arkhangelsky, D., and Imbens, G. W. (2023, 09). Fixed Effects and the Generalized Mundlak Estimator. *The Review of Economic Studies*. Retrieved from <https://doi.org/10.1093/restud/rdad089> doi: 10.1093/restud/rdad089 [4, 13, 14, 28]
- Asquith, B. J., Mast, E., and Reed, D. (2023). Local effects of large new apartment buildings in low-income areas. *The Review of Economics and Statistics*, 105(2), 359–375. [8, 18]
- Ater, I., Shany, A., Ross, B., Turkel, E., and Vasserman, S. (n.d.). Can usage-based pricing reduce traffic congestion? *Work in progress*. [7]
- Athey, S., Blei, D., Donnelly, R., Ruiz, F., and Schmidt, T. (2021). Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data. In *Aea papers and proceedings* (Vol. 108, pp. 64–67). [8]
- Athey, S., Ferguson, B., Gentzkow, M., and Schmidt, T. (2021). Estimating experienced racial segregation in us cities using large-scale gps data. *Proceedings of the National Academy of Sciences*, 118(46), e2026160118. [8]
- Athey, S., Tibshirani, J., and Wager, S. (2019). Generalized random forests. [4, 15, 28, 56]
- Bajari, P., and Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of political economy*, 113(6), 1239–1276. [4, 8, 18, 19, 20, 27]
- Banzhaf, H. S. (2021). Difference-in-differences hedonics. *Journal of Political Economy*, 129(8), 2385–2414. [8, 13, 14]
- Barwick, P. J., Li, S., Waxman, A., Wu, J., and Xia, T. (2024). Efficiency and equity impacts of urban transportation policies with equilibrium sorting. *American Economic Review*, 114(10), 3161–3205. [7, 8]
- Bergé, L. (2018). Efficient estimation of maximum likelihood models with multiple fixed-effects: the R package FENmlm. *CREA Discussion Papers*(13). [5, 26]

- Berry, S., and Pakes, A. (2007). The pure characteristics demand model. *International Economic Review*, 48(4), 1193–1225. [4, 18, 19]
- Bishop, K. C., and Murphy, A. D. (2011). *Incorporating dynamic behavior into the hedonic model*. Unpublished. [18]
- Bishop, K. C., and Murphy, A. D. (2019). Valuing time-varying attributes using the hedonic model: when is a dynamic approach necessary? *Review of Economics and Statistics*, 101(1), 134–145. [18]
- Bordeu, O. (2023). Commuting infrastructure in fragmented cities. [7, 8]
- Borusyak, K., and Hull, P. (2023). Nonrandom exposure to exogenous shocks. *Econometrica*, 91(6), 2155–2185. [16]
- Borusyak, K., Hull, P., and Jaravel, X. (2023). Design-based identification with formula instruments: A review. [16]
- Boston Globe. (2021). *With a new requirement to build housing near train stations, towns brace for density — and drama*. Retrieved from <https://www.bostonglobe.com/2021/12/07/business/with-new-requirement-build-housing-near-train-stations-towns-brace-density-drama/> [35]
- Buchholz, N., Doval, L., Kastl, J., Matějka, F., and Salz, T. (2020). Personalized pricing and the value of time: Evidence from auctioned cab rides. *NBER Working Paper*(w27087). [7]
- Cambridge City Council. (2024). *Housing committee public hearing, may 8th, 2024*. Cambridge, Massachusetts. Retrieved from <https://cambridgema.iqm2.com/citizens/FileOpen.aspx?Type=1&ID=4050&Inline=True> [6]
- Campbell, J. Y., Giglio, S., and Pathak, P. (2011). Forced sales and house prices. *American Economic Review*, 101(5), 2108–2131. [8]
- Cao, Y., Chevalier, J., Williams, K., and Parsley, H. (n.d.). The disparate effects of the retail apocalypse. *Work in progress*. [8]
- Castillo, J. C. (2023). Who benefits from surge pricing? *Available at SSRN 3245533*. [8]
- Chay, K. Y., and Greenstone, M. (2005). Does air quality matter? evidence from the housing market. *Journal of political Economy*, 113(2), 376–424. [8]
- Chen, Q., Syrgkanis, V., and Austern, M. (2022). Debiased machine learning without sample-splitting for stable estimators. *Advances in Neural Information Processing Systems*, 35, 3096–3109. [15]
- Chen, X., Chernozhukov, V., Fernandez-Val, I., Kostyshak, S., and Luo, Y. (2021). Shape-enforcing operators for generic point and interval estimators of functions. *Journal of Machine Learning Research*, 22(220), 1–42. [56]
- Chen, X., and Christensen, T. M. (2018). Optimal sup-norm rates and uniform inference on nonlinear functionals of nonparametric iv regression. *Quantitative Economics*, 9(1), 39–84. [4, 14, 15, 17, 28, 46, 47]
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). *Double/debiased machine learning for treatment and structural parameters*. [4, 14, 15, 28]
- Cook, C. (2022). Heterogeneous preferences for neighborhood amenities: Evidence from gps data. *Available at SSRN 4212524*. [9, 22]
- Cook, C., and Li, P. Z. (2023). Value pricing or lexis lanes? the distributional effects of dynamic tolling. [7, 29, 30]
- Couture, V. (2016). Valuing the consumption benefits of urban density. [9]

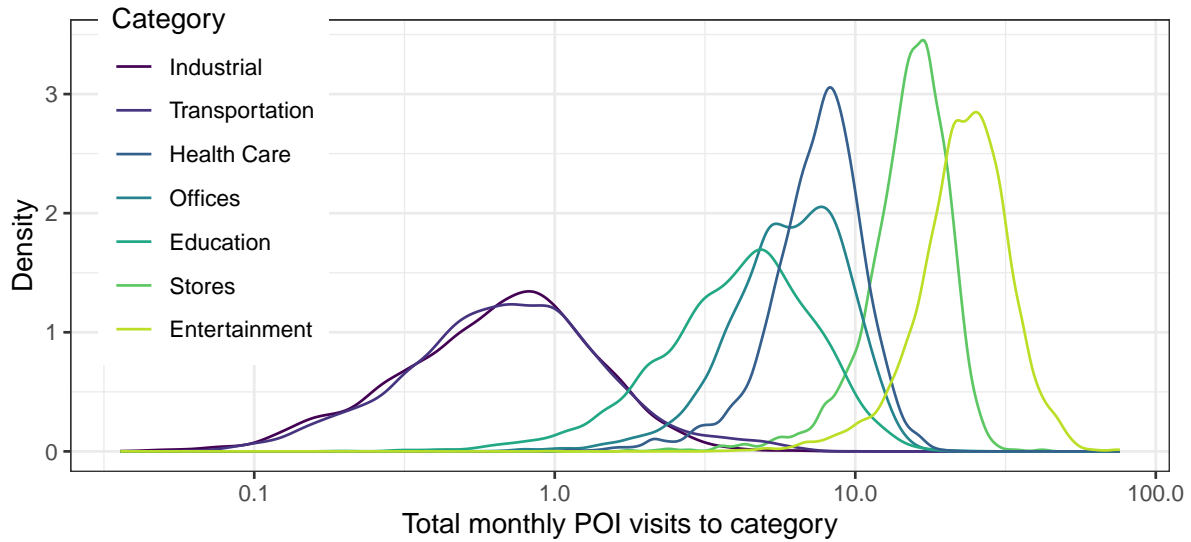
- Couture, V., Duranton, G., and Turner, M. A. (2018). Speed. *Review of Economics and Statistics*, 100(4), 725–739. [6, 30]
- Couture, V., Gaubert, C., Handbury, J., and Hurst, E. (2023). Income growth and the distributional effects of urban spatial sorting. *Review of Economic Studies*. [8]
- Currie, J., and Walker, R. (2011). Traffic congestion and infant health: Evidence from e-zpass. *American Economic Journal: Applied Economics*, 3(1), 65–90. [7]
- Davidoff, T., Pavlov, A., and Somerville, T. (2022). Not in my neighbour’s back yard? laneway homes and neighbours’ property values. *Journal of Urban Economics*, 128, 103405. [8, 18]
- Diamond, R., and McQuade, T. (2019). Who wants affordable housing in their backyard? an equilibrium analysis of low-income property development. *Journal of Political Economy*, 127(3), 1063–1117. [8, 18]
- Dingel, J. I., and Tintelnot, F. (2020). *Spatial economics for granular settings* (Tech. Rep.). National Bureau of Economic Research. [8]
- Duranton, G., and Puga, D. (2023). Urban growth and its aggregate implications. *Econometrica*, 91(6), 2219–2259. [2, 8]
- Durmeyer, I., and Martinez, N. (2022). The welfare consequences of urban traffic regulations. [7]
- Ewing, R., and Dumbaugh, E. (2009). The built environment and traffic safety: a review of empirical evidence. *Journal of Planning Literature*, 23(4), 347–367. [7]
- Glaeser, E., and Gyourko, J. (2018). The economic implications of housing supply. *Journal of economic perspectives*, 32(1), 3–30. [2]
- Goldszmidt, A., List, J. A., Metcalfe, R. D., Muir, I., Smith, V. K., and Wang, J. (2020). *The value of time in the united states: Estimates from nationwide natural field experiments* (Tech. Rep.). National Bureau of Economic Research. [8]
- Gorman, W. M. (1980). A possible procedure for analysing quality differentials in the egg market. *The Review of Economic Studies*, 47(5), 843–856. [4, 18]
- Greenshields, B. D., Bibbins, J. R., Channing, W., and Miller, H. H. (1935). A study of traffic capacity. In *Highway research board proceedings* (Vol. 14, pp. 448–477). [29]
- Greenstone, M., and Gallagher, J. (2008). Does hazardous waste matter? evidence from the housing market and the superfund program. *The Quarterly Journal of Economics*, 123(3), 951–1003. [8]
- Hall, F. L. (1996). Traffic stream characteristics. *Traffic Flow Theory. US Federal Highway Administration*, 36. [6, 29, 30]
- Han, L., Heblich, S., Timmins, C., and Zylberberg, Y. (2024). *Cool cities: The value of urban trees* (Tech. Rep.). National Bureau of Economic Research. [8, 18]
- Herzog, I. (2024a). The city-wide effects of tolling downtown drivers: Evidence from london’s congestion charge. *Journal of Urban Economics*, 144, 103714. [7]
- Herzog, I. (2024b). Microgeographic speed, reliability, and traffic externalities. *Economics of Transportation*. [7]
- Hsieh, C.-T., and Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2), 1–39. [2]
- Kreindler, G. (2024). Peak-hour road congestion pricing: Experimental evidence and equilibrium implications. *Econometrica*, 92(4), 1233–1268. [7]

- Kuminoff, N. V., and Pope, J. C. (2014). Do “capitalization effects” for public goods reveal the public’s willingness to pay? *International Economic Review*, 55(4), 1227–1250. [13]
- Lancaster, K. J. (1966). A new approach to consumer theory. *Journal of political economy*, 74(2), 132–157. [4, 18]
- Li, X. (2022). Do new housing units in your backyard raise your rents? *Journal of Economic Geography*, 22(6), 1309–1352. [8, 18]
- Linden, L., and Rockoff, J. E. (2008). Estimates of the impact of crime risk on property values from megan’s laws. *American Economic Review*, 98(3), 1103–1127. [8]
- Martin, O., and Venugopal, A. (n.d.). Local government speech and affordable housing. [2]
- Mass Live. (2024). *33 towns approved by state under MBTA communities law, eligible for \$15m in grants*. Retrieved from <https://www.masslive.com/news/2024/10/more-than-30-towns-in-full-compliance-with-mbta-communities-law-eligible-for-new-grants.html> [33]
- Massachusetts Executive Office of Housing and Livable Communities. (n.d.). Multi-family zoning requirement for mbta communities. Retrieved from <https://www.mass.gov/info-details/multi-family-zoning-requirement-for-mbta-communities> [35, 51]
- Mast, E. (2023). Jue insight: The effect of new market-rate housing construction on the low-income housing market. *Journal of Urban Economics*, 133, 103383. [2]
- Miyauchi, Y., Nakajima, K., and Redding, S. J. (2021). *The economics of spatial mobility: Theory and evidence using smartphone data* (Tech. Rep.). National Bureau of Economic Research. [8, 22]
- Monte, F., Redding, S. J., and Rossi-Hansberg, E. (2018). Commuting, migration, and local employment elasticities. *American Economic Review*, 108(12), 3855–3890. [8]
- Newey, W. K., and Powell, J. L. (2003). Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5), 1565–1578. [4, 14]
- Oh, R., and Seo, J. (2022). *What causes agglomeration of services? theory and evidence from seoul* (Tech. Rep.). Yale University, mimeo. [8, 22]
- Pennington, K. (2021). Does building new housing cause displacement?: The supply and demand effects of construction in san francisco. [8, 18]
- Pollmann, M. (2020). Causal inference for spatial treatments. *arXiv preprint arXiv:2011.00373*. [8]
- Robinson, P. M. (1988). Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, 931–954. [14, 15]
- Rosaia, N. (n.d.). Competing platforms and transport equilibrium. [8]
- Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy*, 82(1), 34–55. [4, 8, 18, 19]
- Rossi-Hansberg, E., Sarte, P.-D., and Owens III, R. (2010). Housing externalities. *Journal of political Economy*, 118(3), 485–535. [8, 18]
- Severen, C. (2023). Commuting, labor, and housing market effects of mass transportation: Welfare and identification. *Review of Economics and Statistics*, 105(5), 1073–1091. [7, 8]
- Severen, C., and Plantinga, A. J. (2018). Land-use regulations, property values, and rents: Decomposing the effects of the california coastal act. *Journal of Urban Economics*, 107, 65–78. [8]
- Shoup, D. (2018). *Parking and the city*. Routledge. [7]

- Small, K. A., and Chu, X. (2003). Hypercongestion. *Journal of Transport Economics and Policy (JTEP)*, 37(3), 319–352. [6]
- Small, K. A., Winston, C., and Yan, J. (2005). Uncovering the distribution of motorists’ preferences for travel time and reliability. *Econometrica*, 73(4), 1367–1382. [8]
- Su, Y. (2022a). Measuring the value of urban consumption amenities: A time-use approach. *Journal of Urban Economics*, 132, 103495. [5, 9, 29]
- Su, Y. (2022b). The rising value of time and the origin of urban gentrification. *American Economic Journal: Economic Policy*, 14(1), 402–439. [8]
- Swärdh, J.-E., and Genell, A. (2020). Marginal costs of road noise: Estimation, differentiation and policy implications. *Transport Policy*, 88, 24–32. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0967070X17308582> doi: <https://doi.org/10.1016/j.tranpol.2020.01.009> [7]
- Tang, C. K. (2021). The cost of traffic: evidence from the london congestion charge. *Journal of Urban Economics*, 121, 103302. [7]
- Tsivanidis, N. (2024). Evaluating the impact of urban transit infrastructure: Evidence from bogota’s transmilenio. *American Economic Review*. [7, 8]
- Turner, M. A., Haughwout, A., and Van Der Klaauw, W. (2014). Land use regulation and welfare. *Econometrica*, 82(4), 1341–1403. [8]
- USDOT BTS. (2017). *National Household Travel Survey Daily Travel Quick Facts Bureau of Transportation Statistics*. Retrieved 2024-11-04, from <https://www.bts.gov/statistical-products/surveys/national-household-travel-survey-daily-travel-quick-facts> [9]
- Valley Vision. (2023). *The built environment poll*. <https://www.valleyvision.org/resources/the-built-environment-poll-may-2023/>. [2]
- Wang, Y., Wu, H., and Nettleton, D. (2023). Stability of random forests and coverage of random-forest prediction intervals. *Advances in Neural Information Processing Systems*, 36, 31558–31569. [15]
- Wooldridge, J. M. (2021). Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators. *Available at SSRN 3906345*. [14]
- Yang, J., Purevjav, A.-O., and Li, S. (2020). The marginal cost of traffic congestion and road pricing: evidence from a natural experiment in beijing. *American Economic Journal: Economic Policy*, 12(1), 418–453. [6, 7, 29, 30]


A Additional figures and tables

Figure A.1: Trip frequencies by POI category across households



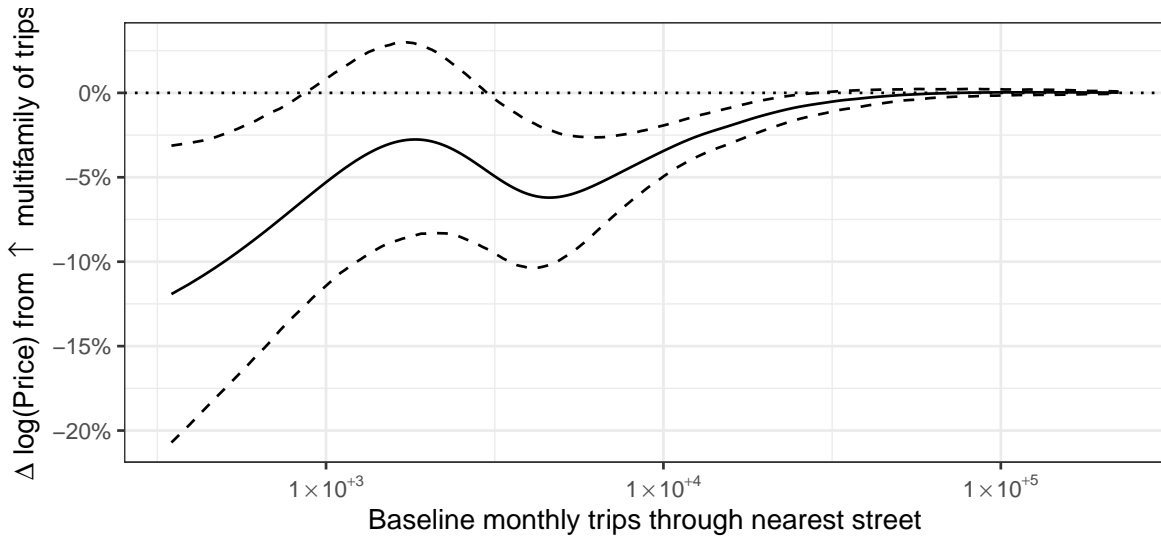
Notes: We present the distributions of monthly visit counts across households to POIs in each of seven aggregate POI categories, which we construct by grouping two-digit NAICS codes. The categories grouping two-digit NAICS codes we construct are Industrial: 11, 22, 23, 31, 32, 33; Transportation: 48, 49; Healthcare: 62, 81, 92; Offices: 51 through 56; Education: 61; Stores: 42, 44, 45; and Leisure: 71, 72.

Figure A.2: Distribution of own-street monthly car trip exposure

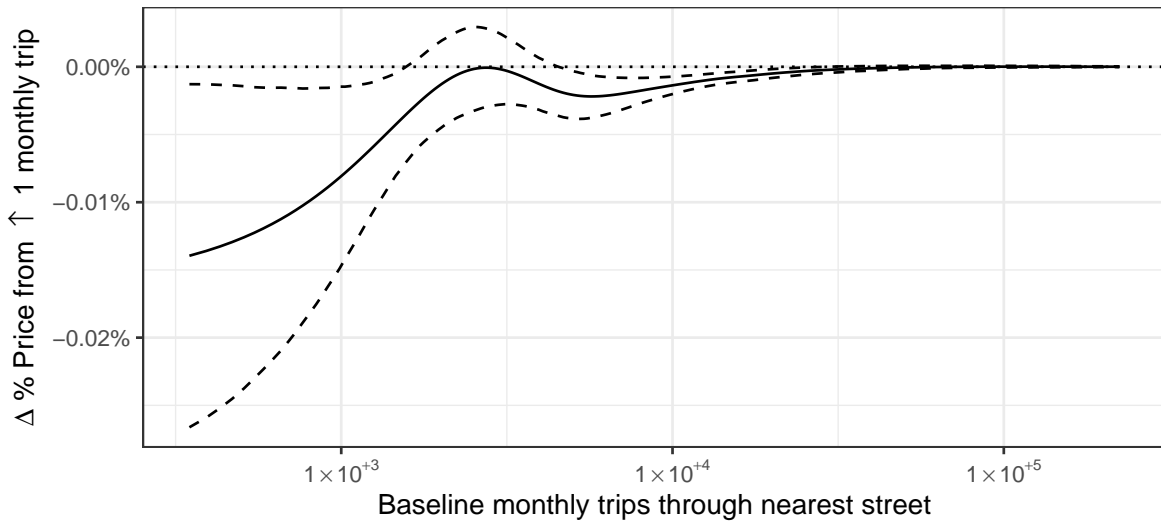


Notes: In this figure, we plot the empirical distribution of monthly car trips down housing units' streets in the Boston CSA.

Figure A.3: Effects of alternative increases in monthly trips on houses' streets



(a) Effects of an additional 3,000 monthly trips



(b) Structural function derivative (effect of one additional trip)

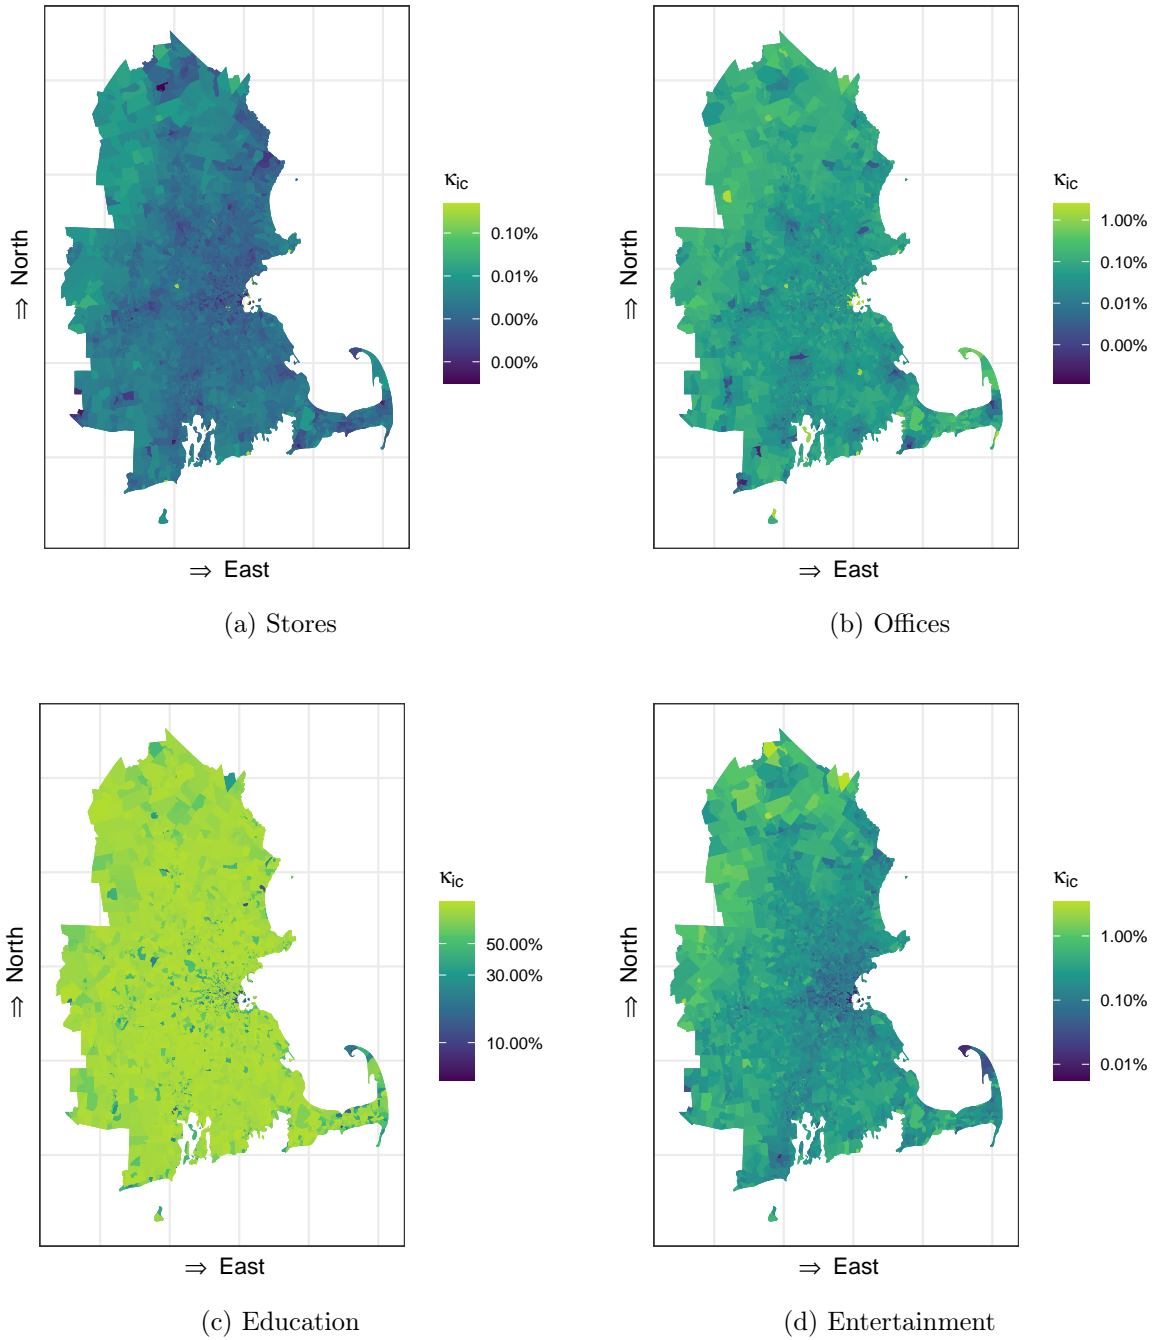
Notes: We present NPIV estimates of two alternative effects of increased nearby street traffic on home prices in 2018, both based on the same estimation procedure described in Section 3.1. In Figure A.3a, we visualize the effects of a 3,000 monthly trip increase in own-street traffic by baseline street traffic. In Figure A.3b, we visualize the derivative of the price function $p_{2018}^{(C)}$, which is essentially equivalent to the effect of adding one additional monthly car trip on a house's street on that house's price. In both figures, dashed lines are 90% uniform confidence bands computed as in X. Chen and Christensen (2018) via a bootstrap clustered at the street level. We instrument for changes in busyness as described in Section 3.1.1. We plot the distribution of nearby monthly trips across in Figure A.2.

Table A.1: Average effects of increases in nearby street traffic by quintile of baseline nearby trips

Baseline trips quintile	Estimate	Std. Err.	95% Conf. Interval	95% Simult. Conf. Int
IQR of monthly trips				
Bottom 20%	-2.91e-01	3.92e-02	[-3.67e-01, -2.16e-01]	[-3.88e-01, -1.95e-01]
20% - 40%	-2.33e-01	3.45e-02	[-3.02e-01, -1.64e-01]	[-3.18e-01, -1.48e-01]
40% - 60%	-1.64e-01	3.01e-02	[-2.22e-01, -1.06e-01]	[-2.38e-01, -8.96e-02]
60% - 80%	-5.01e-02	1.91e-02	[-8.84e-02, -1.19e-02]	[-9.72e-02, -3.02e-03]
Top 20%	-3.75e-04	5.73e-03	[-1.14e-02, 1.06e-02]	[-1.45e-02, 1.38e-02]
3,000 monthly trips				
Bottom 20%	-6.95e-02	2.35e-02	[-1.15e-01, -2.39e-02]	[-1.30e-01, -9.06e-03]
20% - 40%	-4.36e-02	1.71e-02	[-7.72e-02, -1.00e-02]	[-8.75e-02, 2.83e-04]
40% - 60%	-4.61e-02	8.54e-03	[-6.30e-02, -2.92e-02]	[-6.80e-02, -2.41e-02]
60% - 80%	-1.21e-02	3.16e-03	[-1.83e-02, -5.90e-03]	[-2.02e-02, -4.00e-03]
Top 20%	-3.49e-04	7.15e-04	[-1.74e-03, 1.04e-03]	[-2.18e-03, 1.49e-03]
Derivative				
Bottom 20%	-8.98e-05	2.86e-05	[-1.46e-04, -3.40e-05]	[-1.62e-04, -1.73e-05]
20% - 40%	-8.64e-06	8.59e-06	[-2.54e-05, 8.09e-06]	[-3.04e-05, 1.31e-05]
40% - 60%	-1.75e-05	3.81e-06	[-2.49e-05, -1.00e-05]	[-2.71e-05, -7.81e-06]
60% - 80%	-4.60e-06	1.10e-06	[-6.71e-06, -2.48e-06]	[-7.39e-06, -1.81e-06]
Top 20%	-1.37e-07	2.43e-07	[-6.09e-07, 3.34e-07]	[-7.53e-07, 4.78e-07]

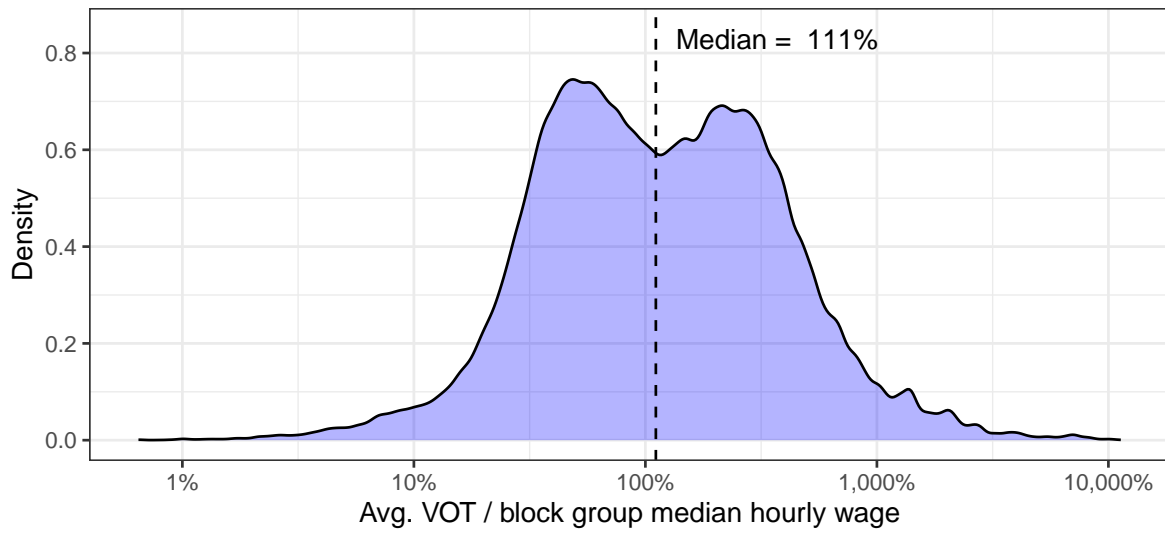
Notes: This table reports average effects in natural log points of different increases in monthly trips on housing units' streets on those housing units' transaction prices. We consider adding the inter-quartile range of trips, 3,000 trips (50 median drivers, each of whom takes 60 trips as seen in Figure 1), and one additional trip, i.e. the derivative. Standard errors, confidence intervals, and simultaneous confidence intervals are computed as in X. Chen and Christensen (2018) via a bootstrap clustered at the street level; the simultaneous confidence bands' simultaneous coverage guarantees hold across the average effect estimates for different baseline street traffic quintiles within each intervention type. We instrument for changes in busyness as described in Section 3.1.1.

Figure A.4: Percent of value-adjusted time κ_{ic} allocated to each POI category



Notes: We present the value-weighted percentages of total POI visit time κ_{ic} allocated to each of four representative categories by block-group in the Boston CSA. These time shares κ_{ic} are calculated as described in Section 4.2.

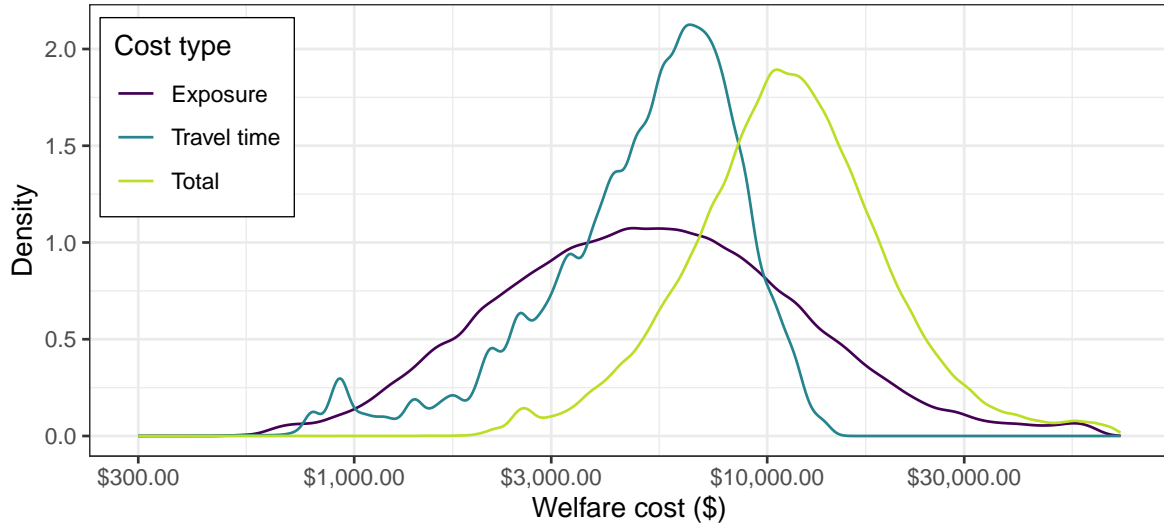
Figure A.5: Distribution of values of time in percentages of hourly wages



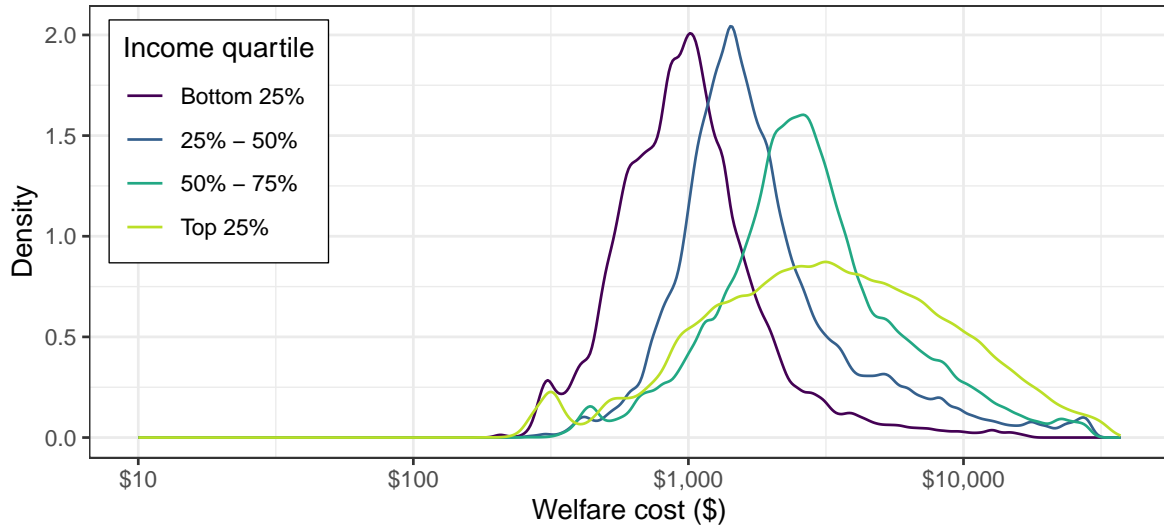
Notes: We present the distribution of households' inferred values of time as percentages of their respective home Census block groups' median household hourly wages, as described in Appendix D.4.

Figure A.6: Distributions of welfare costs per new household

(a) Distributions of total welfare effects across all households by cost type

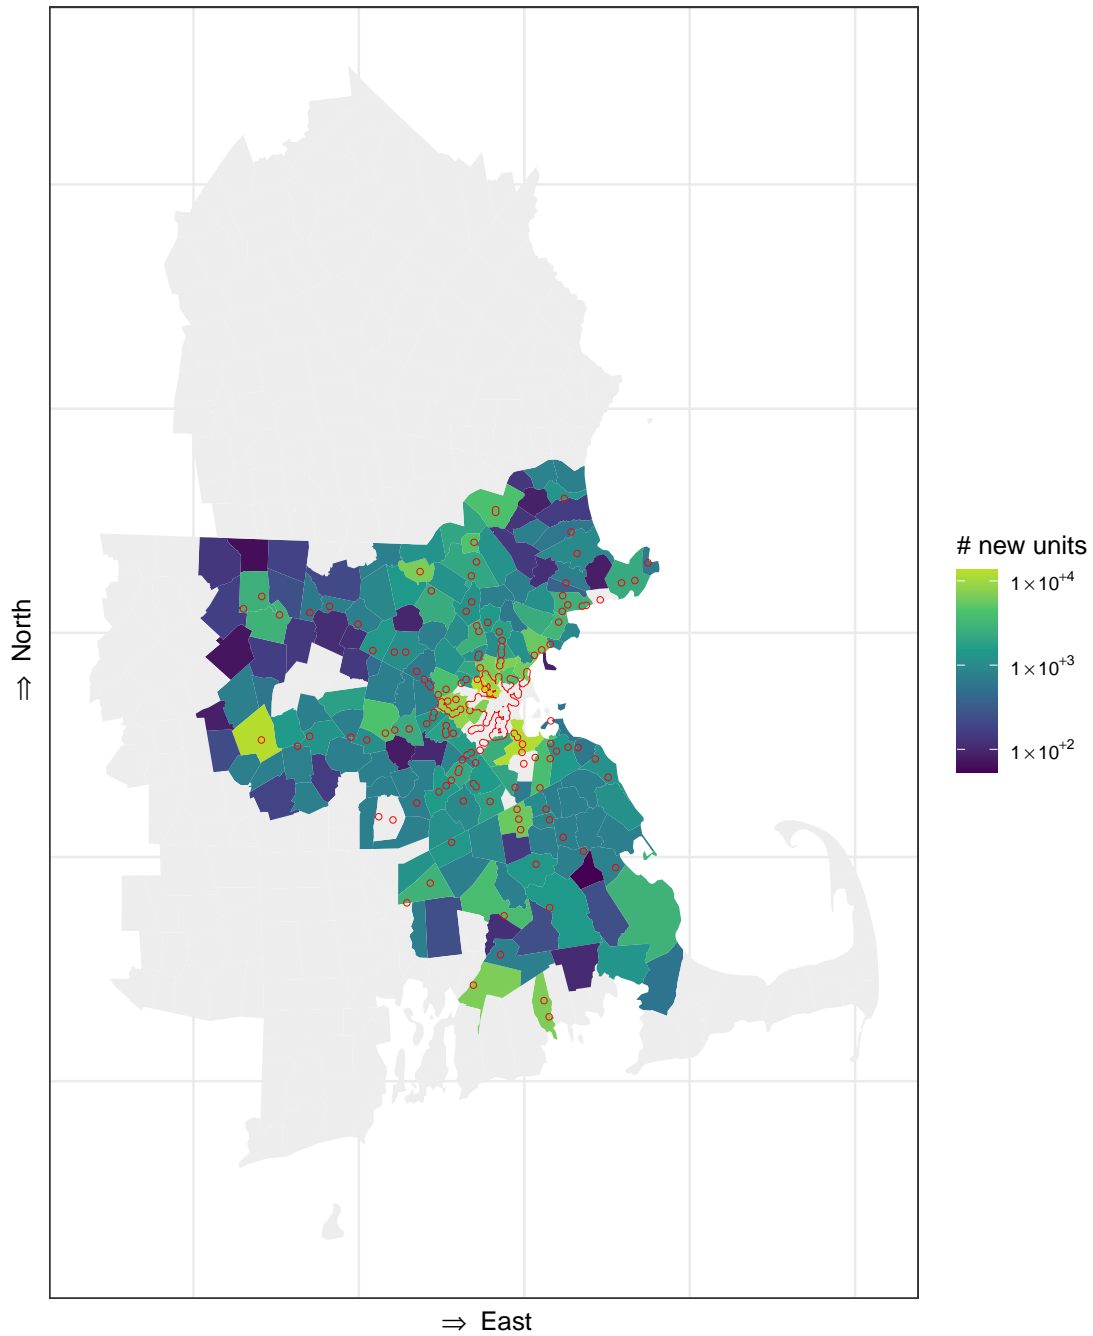


(b) Distributions of total welfare effects by income group



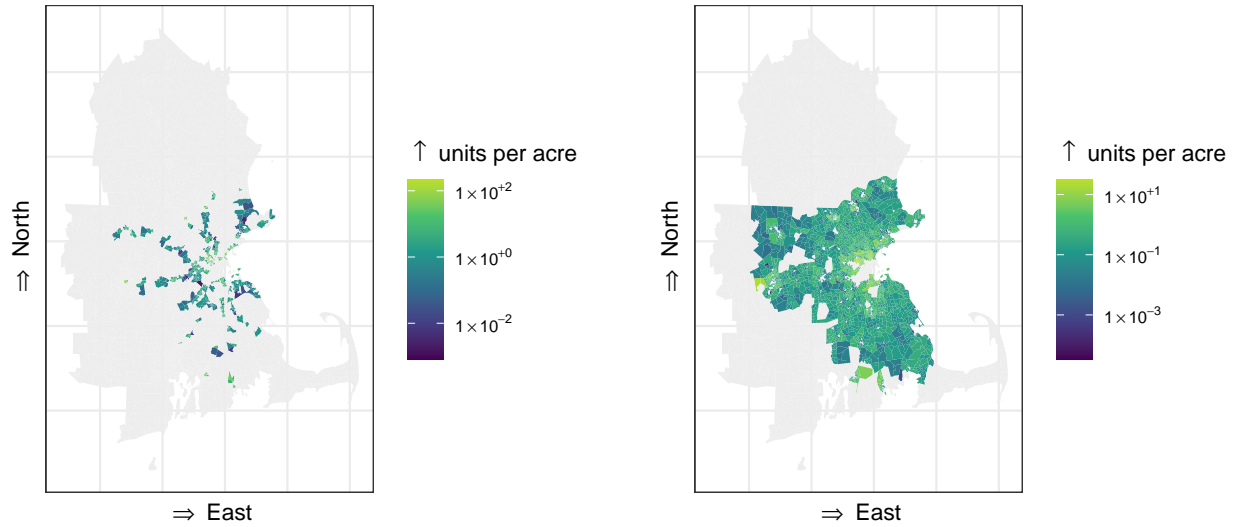
Notes: In Figure A.6a, we plot the distributions of welfare costs by type on incumbent residents from the traffic externalities caused by adding a new household to different locations across the Boston CSA. In Figure A.6b, we plot the distributions of the total welfare costs incurred by households living in Census block groups in different quartiles of the across-block-group distribution of median household incomes. In both figures, each potential new household location is weighted by the total amount of residential land along focal street on which the new housing unit would be sited. We report summary statistics for these distributions in Table 1.

Figure A.7: New unit allocations to MBTA communities and land near public transit stops



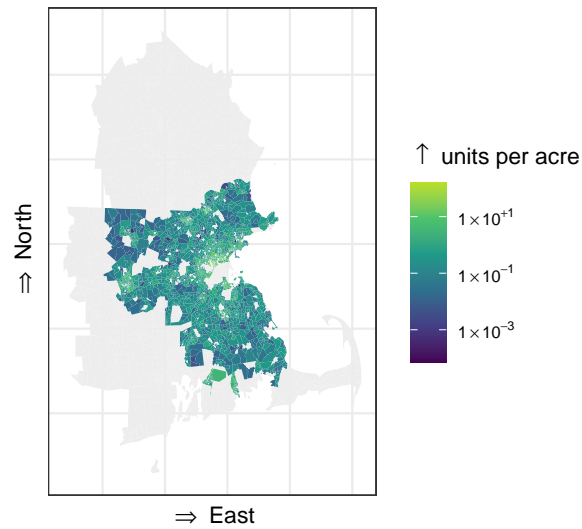
Notes: This figure visualizes each municipality's allocation of new units they must accommodate via land use regulation changes within half a mile of public transit stops under Massachusetts' MBTA Communities law. Areas outlined in red are within half a mile of a public transit stop. Data on municipalities' new unit allocations and eligible land near public transit stops can be found at [Massachusetts Executive Office of Housing and Livable Communities \(n.d.\)](#).

Figure A.8: Minimum density increases by block group for MBTA Communities law variants



(a) Density increases as in MBTA Communities law

(b) Density increases everywhere



(c) Density increases on thoroughfares

Notes: This figure illustrates the increases in minimum density within each Census block group under each of the three policy counterfactuals we consider.

Figure A.9: Traffic costs of MBTA Communities law variants including Boston

Density increase locations	Welfare cost (millions)		
	Total	Exposure	Congestion
All households			
Everywhere	\$4,341	\$2,907	\$1,431
Near transit	\$3,890	\$2,705	\$1,183
On thoroughfares	\$2,844	\$1,512	\$1,329
Income: Bottom 25%			
Everywhere	\$516	\$341	\$175
Near transit	\$602	\$447	\$155
On thoroughfares	\$450	\$284	\$166
Income: 25% to 50%			
Everywhere	\$708	\$481	\$227
Near transit	\$612	\$416	\$196
On thoroughfares	\$491	\$280	\$211
Income: 50% to 75%			
Everywhere	\$1,089	\$713	\$375
Near transit	\$971	\$665	\$304
On thoroughfares	\$702	\$350	\$352
Income: Top 25%			
Everywhere	\$1,999	\$1,250	\$746
Near transit	\$1,648	\$1,039	\$603
On thoroughfares	\$1,272	\$572	\$696

Notes: In this table, we report the total traffic externalities caused by variants of the MBTA Communities law introduced in Section 5.3 where we assume Boston was also allocated housing units in accordance with the MBTA Community law’s guidelines and the rest of the municipalities’ allocations were scaled down proportionally to ensure the same 10% metropolitan-area-wide housing stock increase. Such a change to the policy would require Boston to plan for around 50,500 additional units and would decrease other municipalities’ obligations by 17.5%.

B Details on traffic flow construction

We are not able to observe each household’s routing decision, nor are we able to observe the full distribution of road speeds at the time that a given household is making a trip to a given POI. What we do observe are the ventiles of the speed distributions for each non-residential road in the Boston CSA in our TomTom data, as well as the associated number of cars TomTom used to estimate those speed ventiles. We define 19 traffic states in 2010 and 2018 which are constructed by setting all roads in the Boston CSA to their associated ventiles of observed speeds in that year. We assume that cars drive at the speed limit on residential roads. We compute the fastest driving routes each household would take to reach the POIs they visit under each of these 19 traffic states in 2010 and 2018, amounting to 47 billion unique optimal routes.

While TomTom’s sample sizes are likely smaller than the actual numbers of cars driving through the streets for which we have TomTom data, we will assume that TomTom’s sample of drivers is representative of the population of Boston CSA. As such, TomTom’s sample sizes should be proportional to the true traffic counts on the streets for which TomTom provides road speed data. We use this assumption to estimate the relative frequencies of these aggregate traffic states that best match TomTom’s sample sizes for the road segments that we observe. Our estimated state frequencies need not be interpreted as the true empirical frequencies of each of these traffic distributions, but rather as the frequencies with which drivers choose routes similar to those consistent with each set of state-specific traffic flows. As such, our estimated frequencies account of any misspecification in household’s beliefs over which routes are optimal or suboptimality in households’ routing decisions.

Formally, let \mathcal{L} denote the set of all road segments in the Boston CSA; then for any housing unit $j \in \mathcal{J}$ and POI $\omega \in \Omega$ in the Boston CSA, let $q_{j\omega t}$ denote the number of monthly trips taken from housing unit j to POI ω , and let $\mathcal{R}_{j\omega st} \subset \mathcal{L}$ denote the collection of streets comprising the fastest route from j to ω in traffic state s in period $t \in \{2010, 2018\}$. Then the total monthly traffic flow along street ℓ in period t is given by

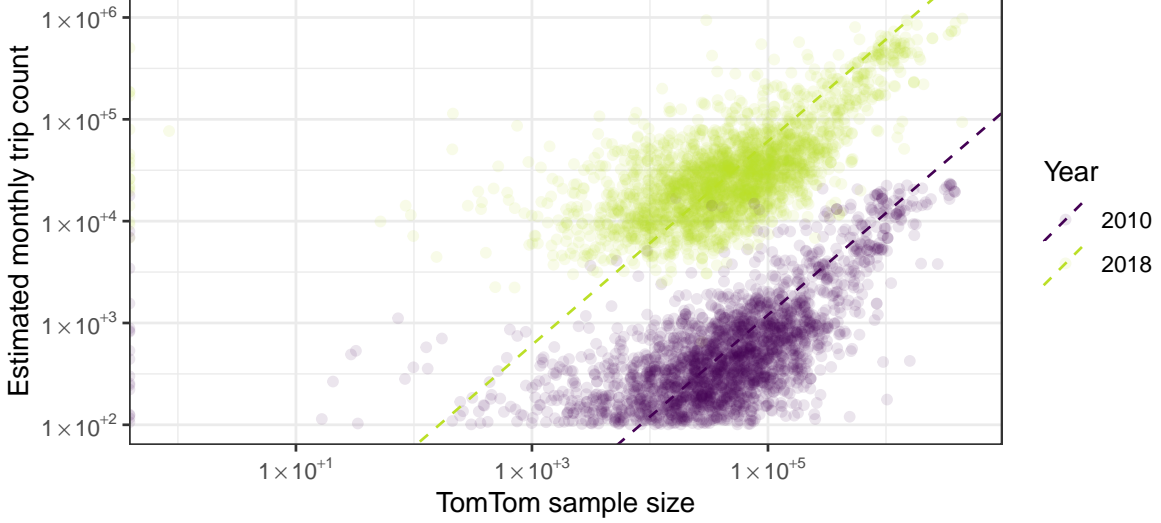
$$C_{\ell ts} := \sum_{j \in \mathcal{J}} \sum_{\omega \in \Omega} \mathbb{1}\{\ell \in \mathcal{R}_{j\omega st}\} q_{j\omega t},$$

and the total monthly traffic flow along street ℓ in year t is given by

$$C_{\ell t} := \sum_{s=1}^S w_{ts} C_{\ell ts}.$$

Given TomTom’s sample sizes $\hat{C}_{\ell t}$ used to construct their speed estimates for a subset of

Figure B.1: Inferred monthly traffic counts and matched TomTom sample sizes



Notes: We visualize a sample of 5,000 estimated monthly traffic flows in 2010 and 2018 through road segments for which we have reported TomTom sample sizes. Dashed lines have unit slopes and intercepts fitted to the data from each year. The R^2 from predicting log-sample sizes using these lines is 73%.

roads $\hat{\mathcal{L}}$, we construct estimates \hat{w}_t of the traffic state probabilities w_t to match TomTom’s sample sizes $\hat{C}_{\ell t}$ via weighted sums of our inferred flows $C_{\ell t s}$ across traffic states up to a multiplicative factor in a squared error sense:

$$(\rho_t, \hat{w}_t) := \arg \min_{\tilde{w} \in \mathbb{R}^S, \tilde{\rho} \in \mathbb{R}} \left\{ \sum_{\ell \in \hat{\mathcal{L}}} \left(\tilde{\rho} \hat{C}_{\ell t} - \sum_{s=1}^S \tilde{w}_s C_{\ell t s} \right)^2 : \tilde{w}_s \geq 0, \sum_{s=1}^S \tilde{w}_s = 1, \tilde{\rho} \geq 0 \right\}.$$

We present a scatter plot of TomTom’s sample sizes and our estimated monthly trip counts on the same streets in Figure B.1.

C Exposure preference details

C.1 Hedonic preference estimation details

We let $\widehat{\log p_{2018}^{(C)}}$ denote our nonparametric estimate of the effect function dictating how the number of monthly trips through a home’s street affects that home’s value in 2018. For houses that transacted in 2018, we can write their “potential” price given a counterfactual quantity of monthly nearby car trips c as follows under the equilibrium price surface model (1):

$$\hat{P}_{j,2018}(c) = P_{j,2018} \exp \left(\widehat{\log p_{2018}^{(C)}}(c) - \widehat{\log p_{2018}^{(C)}}(C_{j,2018}) \right),$$

and we can estimate its derivative with respect to c like so:

$$\frac{\partial \hat{P}_{j,2018}}{\partial c} = P_{j,2018} \exp \left(\widehat{\log p_{2018}^{(C)}}(c) - \widehat{\log p_{2018}^{(C)}}(C_{j,2018}) \right) \frac{\partial \widehat{\log p_{2018}^{(C)}}}{\partial c}.$$

For housing units that did not transact in 2018 and rental apartment units for whom transaction prices at the building level are not easily comparable to home prices, we predict their 2018 prices using a Generalized Random Forest fit to predict the logarithms of 2018 transaction prices minus predicted log-price effects from their nearby street traffic $\widehat{\log p_{2018}^{(C)}}(C_{j,2018})$ (Athey et al., 2019). We use the same property and neighborhood characteristics as we do when estimating our linearly separable controls via Generalized Random Forests in Section 3.1, and we predict rental units' prices as if they were condominiums.⁴⁰ To estimate unobserved housing qualities $\xi_{j,2018}$ for these housing units, we use the $\xi_{j',2018}$ of the nearest transacted unit j' of the same housing type (single-family home or condominium).

Finally, we observe some households choosing to live on streets where our raw price surface derivative estimates are decreasing, which would be inconsistent with those households maximizing their utilities under our hedonic model (3) (Agarwal et al., 2023). To ensure that our estimated price surface is consistent with utility maximization, we apply the post-processing technique described in X. Chen, Chernozhukov, Fernandez-Val, Kostyshak, and Luo (2021) to ensure the potential price functions $\hat{P}_{j,2018}(c)$ are convex and thus consistent with households maximizing utility (Agarwal et al., 2023).

C.2 Using second-order information to bound hedonic utility curvature

In this section, we consider a generalization of the hedonic model of housing utility (3) that allows for concavity in household i 's disutility from exposure to traffic on their street dictated by the parameter $\mu_i \geq 0$:

$$U_{ij}^{(H)} := -\beta_i \frac{(1 + C_{\ell(j)})^{1-\mu_i} - 1}{1 - \mu_i} + U_i^{(A)}(A_j, \eta_{\ell(j)}, \xi_j). \quad (\text{C.1})$$

Note that, by standard properties of isoelastic utility functions, when $\mu_i = 0$, (C.1) reduces to the linear model of disutility from nearby traffic (3) and larger values of μ_i correspond to utility functions with diminishing marginal cost from additional nearby traffic exposure; as $\mu_i \rightarrow 1$, the nearby traffic disutility component of (C.1) approaches logarithmic utility $-\beta_i \ln(1 + C_{\ell(j)})$.

⁴⁰For the single-digit percentages of properties that are missing characteristics, we impute them using the same characteristics of the closest housing unit of the same type (single-family home or condominium).

For simplicity, suppose that for any bundle of other housing characteristics (a, η, ξ) , a housing option with any amount of street traffic c exists within each household's choice set. Then, it suffices to consider household i 's choice of c in isolation, writing the equilibrium price function p as a function only of c and holding their other optimal housing characteristics fixed. As such, we can define household i 's optimal choice C_i^* as follows:

$$C_i^* := \arg \max_{c \geq 0} -\beta_i \frac{(1+c)^{1-\mu_i} - 1}{1-\mu_i} - p(c). \quad (\text{C.2})$$

Note that, so long as p is twice differentiable, C_i^* must satisfy both the first-order condition

$$-\beta_i(1+C_i^*)^{-\mu_i} - \left. \frac{\partial p}{\partial c} \right|_{c=C_i^*} = 0 \implies \beta_i = - \left. \frac{\partial p}{\partial c} \right|_{c=C_i^*} (1+C_i^*)^{\mu_i}, \quad (\text{C.3})$$

and the second-order condition

$$\beta_i \mu_i (1+C_i^*)^{-\mu_i-1} - \left. \frac{\partial^2 p}{\partial c^2} \right|_{c=C_i^*} \leq 0. \quad (\text{C.4})$$

Substituting (C.3) into (C.4) and rearranging, we have that μ_i must satisfy the following inequality:

$$\mu_i \cdot \left[- \left. \frac{\partial p}{\partial c} \right|_{c=C_i^*} \right] \leq (1+C_i^*) \left. \frac{\partial^2 p}{\partial c^2} \right|_{c=C_i^*}.$$

So long as p is strictly decreasing (as should be the case for nearby street traffic), we can then derive the following upper bound on μ_i :

$$\mu_i \leq (1+C_i^*) \left| \left. \frac{\partial p}{\partial c} \right|_{c=C_i^*} \right|^{-1} \left. \frac{\partial^2 p}{\partial c^2} \right|_{c=C_i^*} =: \bar{\mu}_i^{(L)}.$$

If in addition we assume that p is convex in c to ensure that all households' choices are consistent with utility maximization (Agarwal et al., 2023), then the optimization problem (C.2) is a concave maximization problem, so we can strengthen the second-order condition (C.4) to hold globally:

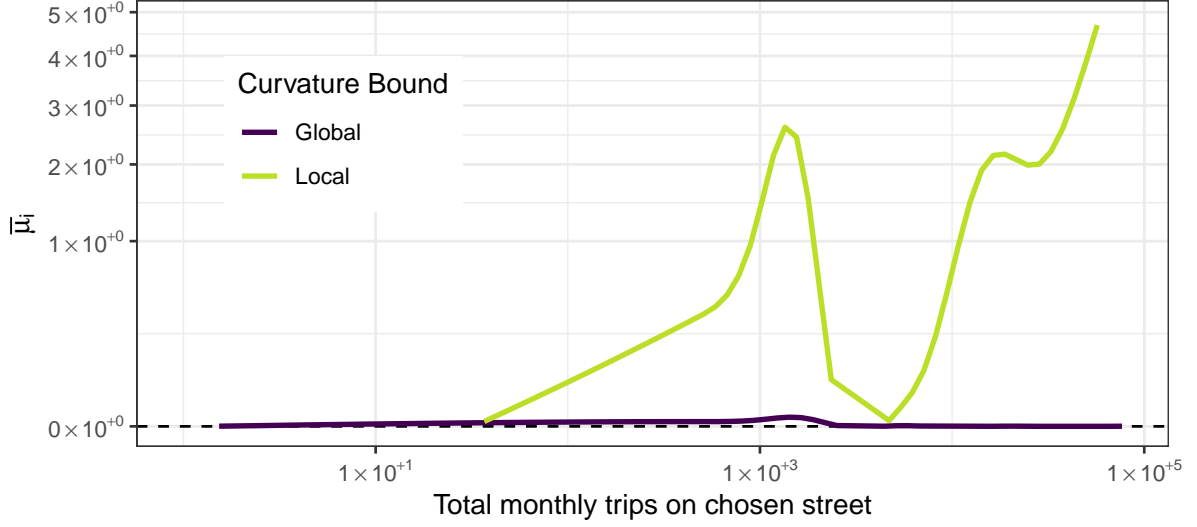
$$\beta_i \mu_i (1+c)^{-\mu_i-1} - \left. \frac{\partial^2 p}{\partial c^2} \right|_{c=c} \leq 0, \quad c \geq 0,$$

which, after solving for β_i as in the first-order condition (C.3) and substituting, implies that

$$\mu_i \left(\frac{1+C_i^*}{1+c} \right)^{\mu_i} \leq (1+c) \left| \left. \frac{\partial p}{\partial c} \right|_{c=C_i^*} \right|^{-1} \left. \frac{\partial^2 p}{\partial c^2} \right|_{c=c}, \quad c \geq 0,$$

where again, we assume that p is strictly decreasing.

Figure C.1: Local and global bounds on utility curvature



Note that since the function $x \mapsto xa^x$ is strictly increasing for all $x \geq 0$ and $a \geq 1$, for each $c \leq C_i^*$, there must exist a $\bar{\mu}_i(c)$ such that

$$\bar{\mu}_i(c) \left(\frac{1 + C_i^*}{1 + c} \right)^{\bar{\mu}_i(c)} = (1 + c) \left| \frac{\partial p}{\partial c} \right|_{c=C_i^*}^{-1} \left| \frac{\partial^2 p}{\partial c^2} \right|_{c=c} \quad \text{and } \mu_i \leq \bar{\mu}_i(c).$$

As it happens, such a $\bar{\mu}_i(c)$ has a closed-form expression in terms of the product log function W , sometimes called the Lambert W function:

$$\bar{\mu}_i(c) = W \left((1 + c) \left| \frac{\partial p}{\partial c} \right|_{c=C_i^*}^{-1} \left| \frac{\partial^2 p}{\partial c^2} \right|_{c=c} [\log(1 + C_i^*) - \log(1 + c)] \right) / [\log(1 + C_i^*) - \log(1 + c)].$$

While W does not have a closed form, it can be evaluated efficiently with a computer. To construct an upper-bound $\bar{\mu}_i^{(G)}$ on μ_i then, we can simply find the minimum value of $\bar{\mu}_i(c)$ over a fine grid of $c \leq C_i^*$.

In Figure C.1, we plot our upper bounds on μ_i obtained using both bounding approaches. Using just local information yields wide bounds on most households' utility curvatures, failing to rule out utility that is significantly more convex in monthly nearby traffic flows than logarithmic utility for many households. However, bounds based on the global property of price surface convexity yields much tighter bounds, ruling essentially any meaningful utility curvature across all households.

D Congestion preference details

D.1 A model of POI visit supply

Here, we describe the details of a monopolistically competitive POI supply model that is consistent with our demand model. We assume category- c , neighborhood- k POIs ω have the same fixed costs F_{ck} of entry and the same constant marginal costs ψ_{ck} of providing a unit of quality to a visitor, and receive the same “price” p_{ck} of providing a visit to a visitor. We put “price” in quotation marks given the discussion in Section 4.1 and conceptualize it more generally as a vertical operationalization of the benefit a POI receives from a visit for the sake of the model. Having entered, POI ω ’s “profit” is then given as follows:

$$\pi_{ck}(v_{ck}(\omega); \omega) := (p_{ck} - \psi_{ck}v_{ck}(\omega)) \mathbb{E}[q_{in(J_i)ck}(v_{ck}(\omega); \omega)].$$

POI ω then chooses $v_{ck}(\omega)$ to maximize “profit”:

$$\begin{aligned} 0 &= \frac{\partial \pi_{ck}(\omega)}{\partial v_{ck}(\omega)} = (p_{ck} - \psi_{ck}v_{ck}(\omega)) \mathbb{E} \left[\frac{\partial q_{in(J_i)ck}(v_{ck}(\omega); \omega)}{\partial v_{ck}(\omega)} \right] - \psi_{ck} \mathbb{E} [q_{in(J_i)ck}(v_{ck}(\omega); \omega)] \\ &\implies (p_{ck} - \psi_{ck}v_{ck}(\omega))(\sigma_c - 1)v_{ck}(\omega)^{\sigma_c - 2} = \psi_{ck}v_{ck}(\omega)^{\sigma_c - 1} \\ &\implies p_{ck}(\sigma_c - 1)v_{ck}(\omega)^{\sigma_c - 2} = \sigma_c \psi_{ck}v_{ck}(\omega)^{\sigma_c - 1} \\ &\implies v_{ck}(\omega) = \frac{p_{ck} \sigma_c - 1}{\psi_{ck} \sigma_c}. \end{aligned} \tag{D.1}$$

By (D.1), $v_{ck}(\omega)$ is constant across POIs ω within category c and neighborhood k . We note that the ratio parameter p_{ck}/ψ_{ck} is identified from (D.1) given v_{ck} and σ_c .

As for POI entry, we assume that category- c POIs enter in neighborhood k until their profits are zero, i.e. until

$$\frac{F_{ck}}{\psi_{ck}} = \left(\frac{p_{ck}}{\psi_{ck}} - v_{ck} \right) \mathbb{E}[q_{in(J_i)ck}(v_{ck})]. \tag{D.2}$$

Since there is not a closed-form solution to (D.2), solving for equilibrium POI entry must be done numerically.

D.2 Derivations of useful POI visit demand quantities

First, by standard results for CES demand,⁴¹ we have that

$$q_{inck}(\omega) = \frac{T_{inc} \sum_{s=1}^S w_s (t_{nks} + d_{ck}(\omega))^{-\sigma_c}}{\sum_{k'=1}^N \int \sum_{s=1}^S w_s (t_{nk's} + d_{ck'}(\omega))^{1-\sigma_c} v_{ck'}(\omega)^{\sigma_c-1} d\mu_{M_{ck'}}(\omega)} v_{ck}(\omega)^{\sigma_c-1}. \quad (\text{D.3})$$

Next, since category- c POIs in k are assumed to be symmetric, qualities $v_{ck}(\omega) = v_{ck}$, dwell times $d_{ck}(\omega) = d_{ck}$, and optimal demands $q_{inck}(\omega) = q_{inck}$, we have that (D.3) becomes

$$q_{inck} = T_{inc} \sum_{s=1}^S w_s (t_{nks} + d_{ck})^{-\sigma_c} v_{ck}^{\sigma_c-1} \cdot \underbrace{\left(\sum_{k'=1}^N M_{ck'} v_{ck'}^{\sigma_c-1} \sum_{s=1}^S w_s (t_{nk's} + d_{ck'})^{1-\sigma_c} \right)^{-1}}_{=I_{nc}^{1-\sigma_c}} \quad (\text{D.4})$$

Again applying within-category-neighborhood POI symmetry, the objective in (5) evaluated at the optimizer (D.4) is given as follows:

$$\begin{aligned} U_{inc} &= \left(\sum_{k=1}^J M_{ck} v_{ck}^{\frac{\sigma_c-1}{\sigma_c}} \left(T_{inc} \sum_{s=1}^S w_s (t_{nks} + d_{ck})^{-\sigma_c} v_{ck}^{\sigma_c-1} I_{nc}^{\sigma_c-1} \right)^{\frac{\sigma_c-1}{\sigma_c}} \right)^{\frac{\sigma_c}{\sigma_c-1}} \\ &= T_{inc} I_{nc}^{\sigma_c-1} \left(\sum_{k=1}^J M_{ck} v_{ck}^{\frac{\sigma_c-1+(\sigma_c-1)^2}{\sigma_c}} \sum_{s=1}^S w_s (t_{nks} + d_{ck})^{-\sigma_c} \right)^{\frac{\sigma_c}{\sigma_c-1}} \\ &= T_{inc} I_{nc}^{\sigma_c-1} I_{nc}^{-\sigma_c} \quad (\sigma_c - 1 + (\sigma_c - 1)^2 = \sigma_c(\sigma_c - 1)) \\ &= T_{inc} I_{nc}^{-1}. \end{aligned}$$

Applying standard results about Cobb-Douglas demand, we have that the optimizers T_{inc} of the objective in (5) are given by

$$T_{inc} = \kappa_{ic} T_i \quad (\text{D.5})$$

Plugging (D.5) into (D.4) yields the expression (7). Finally, plugging (D.5) into the objective in (6) yields the indirect utility expression (8).

D.3 POI visiting model estimation details

Having estimated the within-category travel time elasticities σ_c and the high-dimensional fixed effects ν_{ck} and δ_{cn} , we can back out the visit qualities v_{ck} up to a normalization from

⁴¹See e.g. [these notes](#).

our estimates of σ_c and ν_{ck} :

$$v_{ck} \propto \exp\left(\frac{\nu_{ck} - \log(M_{ck})}{\sigma_c - 1}\right).$$

Given estimates of σ_c and ν_{ck} , if we assume that households' POI visiting preferences are homogenous within each block group, then the following expression yields an estimate of the total time household i who optimally choose to live in neighborhood $n(J_i)$ spends on visits to POIs in category c :

$$\hat{T}_{in(J_i)c} := \frac{1}{N} \sum_{k=1}^N \hat{Q}_{n(J_i)ck} \left(M_{ck} v_{ck}^{\sigma_c - 1} \sum_{s=1}^S w_s (t_{n(J_i)ks} + d_{ck})^{-\sigma_c} I_{n(J_i)c}^{1-\sigma_c} \right)^{-1}.$$

Adding up these estimates $\hat{T}_{in(J_i)c}$ across categories yields household-specific estimates of the total time allocated to POI visits T_i . Finally, by standard Cobb-Douglas utility logic, we can recover the share of time κ_{ic} each household allocates to category- c POI visits:

$$\hat{T}_{in(J_i)c} = \kappa_{ic} \sum_{c=1}^C \hat{T}_{n(J_i)c} \implies \kappa_{ic} = \hat{T}_{n(J_i)c} \left(\sum_{c=1}^C \hat{T}_{n(J_i)c} \right)^{-1}.$$

D.4 Computing households' average values of time

To compute a heuristic estimate of a household's value of time, we first compute the total derivative of household i 's hedonic utility from POI visits across all travel time scenarios:

$$\frac{dU_i^{(T)}}{dt} := -\gamma_i \sum_{c \in \mathcal{C}} \kappa_{ic} \sum_{s=1}^S \frac{\partial \log(I_{nc})}{\partial t_{nks}}.$$

This derivative is a measure of the change in hedonic utility household i would experience were we to increase the travel time for every POI visit they make by an hour every month in perpetuity. We then divide this utility by the total number of monthly trips $Q_{in(J_i)c}$ taken by household i living in neighborhood $n(J_i)$ to construct a heuristic measure of household i 's long-run value of an additional hour of travel time every month in perpetuity on average across the POI visits they make. We then assume households discount the future at an annual rate of 5%, therefore converting these long-run utilities into our estimates of households' monthly values of time plotted in Figure 4 by dividing by $1/(1 - 0.05) = 20$ and then again by 12 (the number of months in a year). To compare these values to households' wages as we do in Figure A.5, we divide our values of time by the median hourly household wages in households' chosen Census block groups, where we compute the median hourly household wage in a block group by dividing the block group's median annual household income by

2,000, a standard estimate of the number of hours a typical person works in a year.