

Who Benefits from Faster Public Transit?

June 12, 2021

Abstract

Lower income commuters are more likely to ride and reside near public transit within cities, but do they also benefit more from faster transit travel? Combining survey data on travel behavior with web-scraped data on counterfactual travel times for millions of trips across 49 large US cities, I estimate a structural model of travel mode and residential location choice. I characterize the heterogeneity across income groups and cities in commuters' willingness to pay for access to faster transit and the expected increase in transit ridership in response to marginal transit improvements. I find that richer transit riders sort more aggressively into the fastest transit routes and are, on average, willing to pay more for faster commutes. Improvements in transit speed are most effective at generating transit ridership and welfare gains where transit is already fast (relative to driving), in cities with a greater share of rail-based transit and where the gains are larger for higher-income commuters. Transit improvements benefit low-income commuters more where transit is relatively slow, in cities with more bus transit, and where the overall marginal gains are small.

Keywords: Travel Mode Choice; Residential Location Choice; Mass Transit; Public Transportation; Income Sorting

JEL classification: R23, R41

1 Introduction

In a rapidly urbanizing world, governments and financial institutions are investing large sums on high-speed inner-city mass transit infrastructure in order to tackle growing road congestion and to reduce carbon emissions.¹ Faster public transit networks are expected to increase transit ridership and reduce the number of vehicles on the road. They are also widely believed to reduce inequality by disproportionately improving mobility and labor market access for the urban poor (Kalachek, 1968). How effective are improvements in travel speed at increasing transit ridership? How are the ridership and welfare gains from faster transit travel distributed across rich and poor commuters? And are the more effective transit improvements also more equitable?

To answer these questions, I develop a discrete choice model of residential location and travel mode choices within cities that reflect heterogenous preferences over travel times by high and low income commuters. To estimate the model, I combine census data on commuting flows and mode choices within US cities with rich web-scraped travel time and route data for millions of counterfactual commuting choices in order to derive the demand for access to faster transit and driving commutes. To the best of my knowledge, this paper is the first to investigate how the demand for faster travel by transit (relative to driving) varies across cities, across commuting routes within cities, and by commuter income. In doing so, I show that improvements in transit speed are most effective at generating overall welfare and transit ridership gains where they benefit higher income commuters relatively more. So, the most effective transit improvements (and the ones likely to be realized) are unlikely to be equitable!

My main findings are as follows. First, I find large differences across cities in commuters' demand for faster public transit commutes. In particular, the marginal willingness to pay (MWTP) for faster transit is significantly higher in cities where transit is already fast (relative to driving) or where a large share of transit usage is via rail transit. For example, the mean MWTP for a one percent increase in commuting speeds among transit riders ranges from \$374 per year in San Francisco CA (and a value of travel time saving of around \$19 per hour) to just \$9 per year in Las Vegas NV.² These

¹For instance, China spent USD 100 billion on rail transit infrastructure in 2017 (OECD, 2019) and opened over 45 subway lines across 25 cities just between August 2016 and December 2017. Hannon et al. [2020] estimates cities and transit providers to undertake at least \$1.4 trillion in new mass transit infrastructure investments by 2025. Over \$100 billion of it will be committed to mass transit in North America.

²Differences across cities in incomes and housing costs play an important role. However, for

differences have important implications for the effectiveness of transit improvements at attracting new transit riders. For instance, a one percent increase in transit speeds throughout San Francisco increases transit ridership by roughly 3.5 percentage points (over 20 percent of baseline transit ridership). In contrast, a one percent increase in transit speeds in Las Vegas increases transit ridership by only 0.2 percentage points (or 6 percent). These city-level patterns mask even larger variation by the location of transit improvements *within* cities. Most notably, the demand for faster transit is significantly higher along commuting routes where transit is already fast relative to driving (such as along rail routes or congested driving routes).

Second, I find that higher income commuters tend to have a higher willingness to pay for faster travel conditional on travel mode choice. Transit improvements attract and benefit lower income commuters more where transit is already slow (relative to driving), as it typically is in most US cities. But transit improvements attract and benefit higher income commuters more where transit is relatively fast. For instance, in New York (the city with the fastest transit speeds in my sample), commuters with annual incomes greater than \$75,000 are 67% more likely to switch to riding transit than commuters with incomes less than \$35,000. In contrast, in Los Angeles (where the transit network is relatively sparse and primarily bus-based), commuters with incomes greater than \$75,000 are only half as likely to switch to transit than commuters with incomes less than \$35,000. Within cities, the income elasticity of demand for faster transit is positive and higher along commuting routes where transit is relatively fast. Together with my first set of results, they imply that the transit improvements most effective at increasing overall welfare and transit ridership are those that benefit and attract higher income commuters relatively more (such as in cities and along popular commuting routes where transit is already fast and driving is slow). This result calls into question the extent to which public transit improvements can be simultaneously efficient and equitable.

Additionally, this paper makes two distinct methodological innovations. The first is a data innovation. While we know anecdotally that travel on public transit is typically slower than on privately-owned vehicles, we have limited understanding of how much faster public transit trips would be on private vehicles (and vice versa) and how they compare across cities and across different parts of the same city. Much of our formal

comparison, the MWTP for a one percent increase in commuting speeds for drivers is \$302 in San Francisco CA (lower than the MWTP of transit riders) and \$17 in Las Vegas NV (much higher than the MWTP of transit riders).

knowledge of travel times to date originates from household travel surveys. But surveys only impart partial information on selected trips which are not directly comparable across mode choices.³ This paper innovates by leveraging a newly emerging source of big data, Google Maps, to predict travel times on both observed and counterfactual trips by each travel mode between the same sets of origin-destination pairs at the same time of day.⁴ The rich variation in travel times across millions of simulated trips allows me to systematically compare transit and driving travel speeds both across and within US cities (and across high- and low-income commutes within cities).

The paper’s second methodological innovation is developing an empirical framework to evaluate the demand for counterfactual transit improvements based off of aggregate cross-sectional data on commuting behavior. To do so, I build on the discrete choice framework developed by McFadden [1978] and extended by Bayer et al. [2004] and Bayer et al. [2007] to recover household preferences for location attributes in the presence of sorting on unobservables. Income sorting into travel modes and neighborhoods necessarily induces correlations between location attributes endogenous to incomes and other unobservable (and observable) location attributes.⁵ My empirical framework gets around such endogeneity concerns by allowing choices to condition on unobservable attributes of travel modes and residential neighborhoods. In addition, my paper extends on this class of residential sorting models by conditioning out mean preferences across income groups over the unobservable attributes of travel modes (thus essentially controlling for the income sorting). Then, preferences over commuting speeds are identified from the residual variation in individual workers’ commutes to their given work locations within the city.

Identifying preferences off of net variation in commuting speeds instead of proximity to transit (as is common among studies of inner-city transit) makes a big difference to the estimated distributional gains from faster transit: because while poorer commuters tend to reside closer to transit stops in typically high-density neighborhoods, I show that richer commuters are the ones who enjoy the fastest transit commutes within cities. Who benefits more from improvements in transit speed (and how much) depends on

³Self-reported travel times are also subject to recall bias, anchoring and related measurement concerns.

⁴Google Maps exploits historical and real-time data from tracking the movement of smartphones combined with information on transit schedules to predict travel times that have been shown to credibly capture variation in travel times from real driving trips (Akbar et al., 2018).

⁵For instance, higher-income neighborhoods may be higher-quality or have higher travel speeds because of better-funded local amenities and infrastructure.

how fast transit is (already) relative to driving.

My investigation has important implications for three broad groups of literature. First, the paper’s findings help us better understand public transit’s role in neighborhood gentrification and inequality. Because poorer commuters have been shown to be more likely to ride transit and reside near new transit stops (Baum-Snow and Kahn, 2000, 2005, Pathak et al., 2017), public transportation in the US is frequently portrayed as an inferior good and a poverty magnet (Glaeser et al., 2008). This paper shows that transit is indeed more likely to attract poorer commuters where it is slow relative to driving (as it is on most commutes). But as we make transit faster, it becomes relatively more attractive to the rich (and a normal good). This result is consistent with recent individual case studies of high-speed transit expansions, which often find incomes to have gone up in newly transit accessible neighborhoods (Heilmann, 2018) and richer commuters to have benefited just as much or more (Tsivanidis, 2019). It may be that realized high-speed transit expansions often attract the rich more because planners are focusing on efficiency rather than equity (as suggested by my results), and transit expansions that are more attractive to the poor would also be less effective overall (as in Gaduh et al., 2020).

Second, this paper informs us of the value of travel time savings on public transit. Papers comparing the effect of different public transit expansions have overwhelmingly focused on proximity to transit stations or distance along transit routes assuming anecdotal or constant speeds (Kahn, 2007, Glaeser et al., 2008, Gu et al., 2019, Pathak et al., 2017 to name a few). In contrast, my data allows me to directly estimate preferences for faster transit commute (instead of proximity to transit). There is also a large literature on measuring people’s opportunity cost of time spent traveling and using it to inform transportation policies at the intensive margin, such as for congestion pricing (Small, 2012, Bucholz et al., 2020, Goldszmidt et al., 2020). While the value of travel time savings (VTTS) has been extensively studied based on driving trips, this paper tells us about the VTTS on public transit and how it compares to driving across income brackets and across cities with different transit networks. The distinction proves important because I estimate VTTS among mass transit riders that are, on average, half the VTTS among drivers.⁶

⁶Craig [2019] also uses residential location and travel mode choices to estimate the value of commuting time in Vancouver, but their model does not distinguish the value of time by each mode of travel. Also, their variation in travel time is based on reported transit schedules and survey-reported driving times.

Third, this paper contributes to the literature quantifying the gains from investing in inner-city mass transit infrastructure. A growing number of case studies of individual mass transit expansions investigate transit’s general equilibrium effects on the spatial distribution of economic activity within cities (Heblich et al., 2020, Severen, 2019, Tsivanidis, 2019, Warnes, 2020). In order to precisely estimate the heterogeneity in preferences for faster transit commutes, I deviate from these quantitative spatial models by foregoing much of the restrictions on preferences imposed by their model structure. Instead, I use a more flexible utility specification that allows me to more precisely identify the heterogeneity in preferences across income groups and across space.⁷ Understanding this heterogeneity is key to be able to generalize case studies to inform policy making. For instance, how informative are model predictions for one city about potential transit improvements in another? I show that the demand for faster commutes varies widely but systematically across cities and across locations within cities.

This paper focuses only on the direct travel time gains from mass transit improvements, which Tsivanidis [2019] found to have accounted for 60-80% of the total welfare gains in general equilibrium from expanding Bus Rapid Transit in Bogotá. There are also studies that explore mass transit’s implications for population decentralization (Gonzalez-Navarro and Turner, 2018, Lin, 2017), income segregation (Akbar, 2020), congestion (Anderson, 2014, Gu et al., 2019), car ownership (Mulalic et al., 2020), air pollution (Gendron-Carrier et al., 2020), property values (Bowes and Ihlanfeldt, 2001, Cervero and Kang, 2011, Gupta et al., 2020), labor market informality (Zárate, 2020), gender inequality (Kwon, 2020, Kondylis et al., 2020) and long-term growth of cities (Heblich et al., 2020) among other things.

The rest of this paper is organized as follows. Section 2 describes the available data on observed commutes and the data estimation process for counterfactual commutes. Section 3 documents differences in transit ridership and transit travel times relative to driving (both across cities and within cities). It also documents differences across income groups in their access to high-speed transit and driving commuting routes. Section 4 presents a model of travel mode and residential location choices within cities and an estimation strategy to identify the demand for faster transit commutes. Section 5 presents the estimated preferences in terms of commuters’ willingness to pay for faster travel and characterizes the heterogeneous gains in transit ridership and welfare from

⁷In doing so, I also forego the ability (of these models) to simulate mass transit’s longer term implications for urban residents beyond the immediate gains from faster travel.

marginal improvements in transit speeds across cities. Section 6 concludes.

2 Data

This paper studies the residential neighborhood and travel mode choices of commuters in the 2006-10 American Community Survey (ACS). A 'city' in this paper is a metropolitan area (CBSA) and I focus on the 49 metropolitan areas with population over 300,000 where at least 2% of the commutes are by public transportation.

2.1 Commuting flows

Data on the flow of commuters between each pair of residence and work census tracts comes from the 2006-10 Census Transportation Planning Package (CTPP), which are aggregations of the ACS microdata for the corresponding years. I use the breakdown of the population of commuters by their household income bracket and their means of transportation to work.⁸ Household income brackets are fixed for all metropolitan areas at (1) under \$35,000, (2) \$35,000-50,000, (3) \$50,000-75,000 and (4) over \$75,000. I restrict my analysis to workers over 16 years old who commute to work within the extent of my CBSAs and who either drive, ride public transportation or walk to work.⁹ For the rest of the paper, I use the term 'transit' to refer exclusively to public transportation.

Across all cities, my sample covers roughly 61 million commutes across 2 million observed residence-work tract pairs. The vast majority of commuters in each city choose to drive. The fraction of commutes by transit is 4% in the median CBSA in my sample and is as high as 31% (in New York, NY). The fraction of commutes by 'walking' is around 3% in the median CBSA and as high as 11% (in Boulder, CO).

In addition to the aggregate counts of commuting flows, my analysis relies on housing expenditure data on individual workers from the 5% sample of microdata from IPUMS (Ruggles et al., 2019) and aggregate demographic data on residential census tracts and block groups (with more detailed breakdown on household incomes and

⁸CTPP data for more recent years do not include these tabulations for the interaction of household income and means of transportation. Thus, my analysis is limited to ACS years 2006-10.

⁹Driving pools together both those who ride their own vehicle and those who carpool with others on privately owned vehicles. Unfortunately, walking includes bicycling as the ACS lumped together counts of commuters who walk to work with those who bike.

choices than the CTPP data) from the National Historic Geographic Information System (NHGIS). I use population-weighted centroids and crosswalks between geographies from the Missouri Census Data Center. To measure housing prices, I use standardized property prices on single-family parcels at the census tract level from Davis et al. [2020].

2.2 Travel times

The analysis in this paper relies on knowing the travel times faced by each worker from their observed residential locations and travel modes as well as from (their unchosen) alternative locations and modes. To construct these counterfactual travel times, I simulate a series of trips on Google Maps by driving, transit, and walking from every block group in the CBSA at exactly the same time of the same day. These include trips to nearby popular shopping malls, restaurants, schools, pharmacies and 15 other destination types from Google’s directory of “place types” (the exact trip destinations depend both on the destination’s popularity as a Google search result as well as on its proximity to the trip origin). I also include trips to the 5 most popular work destinations in the residential county as well as to the 5 most popular work destinations in the residential census tract (based on commuting flows observed in the CTPP data). Then I use the spatial variation in Google’s travel time predictions on these trips to impute travel times on all possible counterfactual commuting trips.

Google’s travel time predictions on trips by driving and walking are based on their historical data on smartphone movements.¹⁰ In contrast, travel time predictions on trips by transit are based on schedules shared by local transit authorities (sometimes in real-time) and the open-source General Transit Feed Specification (GTFS). These transit travel times include waits between transfers as well as time spent walking to transit stops. For trips with no viable transit routes nearby, Google returns the predicted walking times. Since transit travel times are sensitive to the timing of the Google Maps query and the departure time (which are not planned relative to the transit schedules unlike most real transit trips), I search each trip at five different times of the day and consider a weighted average of the travel times in subsequent analyses.¹¹ I only do so for transit trips as the driving and walking travel times returned by Google are

¹⁰Google also makes real-time travel time predictions but they are more susceptible to idiosyncratic shocks and the timing of the data collection.

¹¹The weights are proportional to the hourly frequency of trips (by trip purpose) in the 2017 NHTS.

already historical weighted averages across time of day and days. Appendix Section A.1 includes additional details on identifying trip destinations and querying trips on Google Maps.

Importantly, I need to predict the counterfactual commuting times faced by individual workers from their observed work tracts to each residential tract within the CBSA by each travel mode. There are 38 million residence-work location combinations faced by commuters across my 49 cities. Since the full matrix of possible trips between these location pairs by each travel mode is too large (and expensive) to query individually on Google Maps, I rely on an alternative approach that proceeds in three steps. First, I identify the shortest routes between all trip origin-destination pairs (including for the non-commuting trips queried on Google Maps) along major road networks downloaded from OpenStreetMap (OSM) and compute the overlap between these routes and the city’s tract boundaries. Second, using the 14 million trips for which I have travel times from Google Maps, I estimate the average speed of travel through each tract by each mode.¹² Third, I use the estimated mode-tract-specific speeds and the overlaps between tracts and routes to predict travel times on the remaining (commuting) trips. I repeat the three steps separately for each CBSA.

More precisely, let τ_{cqm} denote the travel time on trip q in CBSA c using travel mode m . I can decompose it into a sum of travel times through each overlapping tract on its route:

$$\tau_{cqm} = \sum_{k \in K_c} l_{ckq} / s_{ckm} \quad (2.1)$$

where l_{ckq} is the trip length overlapping tract k , s_{ckm} is the mode-specific travel speed through tract k , and K_c is the set of census tracts within a convex hull of the CBSA’s geographic extent. To determine the overlap l_{ckq} between trips and tracts, I compute each trip’s shortest route along the network of non-residential streets and intersect it with all tract boundaries. Then, using the set of trips for which I also have total trip travel times τ_{cqm} from Google Maps, I estimate travel speeds s_{ckm} from (2.1) using an OLS regression of the trip travel times on the trip lengths overlapping each tract (with coefficients $1/s_{ckm}$). I run these regressions separately for each CBSA and travel mode. Then I plug the estimated speeds into (2.1) to predict travel times on the commuting

¹²In this version of the paper, the Google Maps travel times to non-residential amenities (such as restaurants, shopping malls and parks) only serve to help me predict travel times on commutes since this study focuses only on commuting trips. An extension (in progress) investigates worker preferences on both commuting trips and trips to amenities.

trips that did not get queried on Google Maps.¹³ See Appendix Section A.2 for further details on estimating tract speeds and commuting times.

As a quality check, I use the estimated speeds to also predict travel times for a randomly selected test sample of trips for which I already have Google travel times but which I do not use in the speed estimation. The predicted travel times are strongly correlated with the Google travel times: for the median city, the correlation for driving and walking travel times are 95% and 97% respectively. The median correlation between transit travel times is slightly lower (but still reasonably high) at 84%.

3 Travel Speeds, Mode Choices and Incomes

Henceforth in the paper, trip “distance” refers to the length of the shortest OSM route and the (average) trip “speed” is this shortest route distance divided by the predicted travel time. Note that trip distances are not necessarily the traveled road distances and do not vary with travel modes. Accordingly, trips speeds measure both how fast one moves along a route as well as the directness of the travel route (relative to the shortest street route). For example, higher transit speeds on a trip may correspond to either a more direct transit connection (such as one with fewer detours or less time spent walking and waiting along the way) or a faster transit route (such as one with fewer stops in between or by subway instead of bus).¹⁴ Similarly, driving speeds reflect both the directness of chosen driving routes (relative to the shortest route along major arteries) and how fast traffic flows along these routes. This definition of speed is arguably the more policy-relevant measure of interest: how well the transit or driving network connects locations within cities (rather than how fast vehicles move).

¹³An advantage of using commuting travel times predicted from these tract-level speeds instead of travel times directly returned by Google Maps on a trip between the (centroids of) tracts is that they are less sensitive to how one chooses the precise locations of the trip origins and destinations within tracts. It pertains even more to transit travel times because transit routes can be sparse and walking times to and from transit stops can vary significantly depending on where the trip starts and ends. The tract level speeds smooth out this variation within tracts. So, while the predicted commuting times may not be the best predictor of actual travel times between the tract centroids, they may be more representative of average travel times between the tracts.

¹⁴The transit speeds do not include scheduling costs related to when to start the trip. For instance, Google Maps may ask the traveler to start their trip at a particular departure time to have them coincide with the arrival of a bus or train at the designated stop. The difference between the scheduled departure time and the time the trip is queried is not included in the travel times. In subsequent analyses in Section 4 onwards, this schedule cost is covered by travel mode-origin fixed effects. Note, however, that wait times between transit transfers are included in the travel times.

Table 1 shows mean travel distances, times and speeds across all commuting trips in my sample conditional on each commuter choosing their observed residence, work location and travel mode. On average, drivers reside farther away from work than transit riders, who reside farther away than walkers. Unsurprisingly, driving commutes are also faster than transit commutes, which are faster than walking commutes.¹⁵ But regardless of mode choice, higher income commuters tend to commute both longer (in distance and time) and faster than lower income commuters.

Table 1: Mean distances, times and speeds on observed commutes

	Travel Mode	All commuters	<\$35k	\$35k-\$50k	\$50k-\$75k	>\$75k
Distance (in km)	driving	23.2	20.2	21.4	22.6	24.3
	transit	22.6	15.2	17.2	19.4	27.2
	walking	8.2	7.7	7.9	8.0	8.7
Travel time (in min)	driving	22.7	20.3	21.2	22.1	23.5
	transit	74.5	58.0	62.6	67.4	84.8
	walking	87.6	81.4	84.6	85.0	93.8
Speed (in km/h)	driving	55.2	52.6	53.8	54.9	55.9
	transit	17.1	14.8	15.6	16.4	18.5
	walking	4.9	4.8	4.9	4.9	4.9

Note: Means are over all observed one-way commutes, i.e. conditional on each commuter choosing their observed residence, work location and travel mode. Table pools together all cities in my sample.

¹⁵Recall that walking commuters include bicyclists but the mean travel times and speeds are based on Google Maps' predictions for walking. So, the average walker does not actually spend 88 minutes on commute - many of the longer trips presumably happen via bicycle.

Table 2: **Ranking of cities by mean commuting speeds on transit**

Rank	City	Transit speed (in km/h)	Ratio of transit to driving speed	% commuters riding transit	Rail share of transit riders
1	New York, NY	20.2	0.35	30.7%	86.7%
2	San Francisco, CA	18.9	0.31	15.5%	51.6%
3	Seattle, WA	17.8	0.30	8.6%	5.2%
4	Chicago, IL	17.7	0.28	11.9%	69.8%
5	Philadelphia, PA	17.1	0.30	9.7%	46.0%
-	-	-	-	-	-
10	Atlanta, GA	15.7	0.21	3.5%	30.8%
-	-	-	-	-	-
15	Minneapolis, MN	15.1	0.20	4.8%	5.8%
-	-	-	-	-	-
26	San Diego, CA	13.9	0.24	3.5%	11.0%
-	-	-	-	-	-
37	Austin, TX	12.8	0.21	2.8%	1.1%
-	-	-	-	-	-
49	Vallejo-Fairfield, CA	8.8	0.18	2.7%	33.3%

Note: Speeds are relative to shortest road distance (not necessarily the travel distance). Speeds and ratios of travel times are means across all trips between observed work-residence location pairs (unconditional on travel mode choice) ignoring the top and bottom 5% of outliers. Rail share is the fraction of transit commutes via rail transit in the city.

Travel speeds also vary greatly across cities. Table 2 lists the fastest cities by their mean transit speeds across all observed residence-work location pairs *unconditional on travel mode choice*.¹⁶ The table also reports the means of transit speeds as a function of driving speeds on corresponding trips. On average, transit is slower than driving everywhere. In the fastest transit cities, driving is roughly three times faster than riding transit. In the slowest transit cities, driving is roughly five times faster. Cities with relatively faster transit have a higher share of commuters who ride transit. These cities are also likely to have a higher share of their transit commutes using rail-based transit as opposed to road-based transit (such as buses). But there are exceptions, consistent with the speeds reflecting both how fast commuters move along their transit routes as well as how well-connected transit routes are. Most notably, Seattle WA has the third highest average transit speed but 95% of its transit commutes are by bus, whereas Vallejo-Fairfield CA has the lowest average transit speed but a third of its fewer transit riders are more likely to use commuter rail.

¹⁶Appendix Table A.1 shows a complete ranking of all cities by transit speed.





Figure 3.1: **Share of commuters riding transit as a function of travel speeds.** Figure pools together commutes across all cities. Speeds are standardized (to mean 0 and std. dev. 1) across trips between observed work-residence pairs within each CBSA. Trips at the top and bottom percentiles of speeds are ignored. White spaces in the graph correspond to 0.1-by-0.1 cells with fewer than 20 commutes.

Within cities, the likelihood of riding transit depends on the speed of transit relative to driving. Figure 3.1 plots commuters' probability of riding transit (on the z-axis) against commuting speeds by driving and transit between their observed residence-work location pair (on the x- and y-axes). I standardize the transit and driving speeds within each city so that we are comparing mode choices across locations within (and not across) cities. Conditional on driving speed, transit ridership is higher on trips (i.e., residence-work location pairs) where transit is relatively fast. Also, conditional on transit speed, transit ridership is higher on trips where driving is slow. As such, transit and driving are substitutable alternatives: commuters choose more of one when the price of travel (in terms of inverse travel speeds) on the other is higher.



(a) High-income share of **transit riders** by their commuting speed (b) High-income share of **drivers** by their commuting speed

Figure 3.2: Share of commuters who are high-income by (standardized) speed of travel. The figures pool together commutes across all CBSAs. Horizontal axis depicts travel speeds on chosen mode standardized (to mean 0 and std. dev. 1) across all observed commutes on the same mode within each CBSA. Trips at the top and bottom percentiles of speeds are ignored. Speed, in this context, is the shortest road distance divided by travel time. Confidence intervals are in grey.

The income composition of transit riders also varies systematically over travel speeds. As shown in Figure 3.2a, higher income transit riders are more likely to enjoy faster transit commutes within a city.¹⁷ A similar pattern can be observed for drivers. Figure 3.2b shows that higher income drivers are more likely to enjoy the fastest driving commutes (but the mean differences are smaller than among transit riders). These patterns could be due to an income-elastic preference for faster commutes that is missed if we focus only on travel times instead of speeds. As explored further in Appendix Section A.5, when commuters travel faster, they also commute longer. And, as seen in Table 1, higher income commuters have higher average travel times (and distances) on their chosen travel mode despite higher travel speeds.

¹⁷The observed relationship is robust to alternative high income cutoffs too besides \$50,000.




Figure 3.3: Mean incomes of commuters as a function of driving and transit speeds. Figure pools together commutes across all cities. Household incomes are means across commutes of medians of income brackets (based on micro-data). Speeds are standardized (to mean 0 and std. dev. 1) across trips between observed work-residence pairs within each CBSA. Trips at the top and bottom percentiles of speeds are ignored. White spaces in the graph correspond to 0.1-by-0.1 cells with fewer than 20 commutes.

Having said that, unconditional on mode choice, higher income commuters are more likely to sort into commutes where driving is fast. Figure 3.3 shows average commuter incomes by driving and transit speeds between their work and residence (with redder shades now depicting higher mean commuter incomes). Given both driving and transit commutes appear to be normal goods and driving is typically faster (cheaper in time) than transit, it is unsurprising that average commuter incomes are higher where driving is faster. In a few cases, incomes are also high where transit is fast and driving is not. Income sorting across locations appear to be reflective of heterogeneous preferences over commuting speeds and travel modes, but the sorting could also be driven by preferences over other spatially correlated features. The following section proposes an empirical framework to isolate the extent to which the observed income sorting into work-residence location pairs and travel modes can inform us about the gains from a faster transit network.

4 A model of travel mode and residential location choice

Suppose each city is composed of a fixed population of heterogeneous workers, a set of residential neighborhoods n and work locations j , and three modes of travel $m \in M = \{\text{driving, transit, walking}\}$. Workers classify under one of four income groups y , each with a fixed population in the city and a different distribution of jobs across work locations. Each worker i is exogenously assigned a work location and an income w_i , and choose their residential neighborhood and mode of travel to maximize their gains from shorter commutes given heterogeneous preferences over mode and neighborhood characteristics. The rest of this section characterizes (and parameterizes) the worker's decision problem and outlines a strategy to estimate the preference parameters from available data.

4.1 Specification

Work locations determine the set of commuting times workers face to each residential neighborhood by each travel mode, and workers from different income groups may have different preferences over these commuting times. For instance, if higher income commuters have a higher opportunity cost of time, they are likely to have a stronger preference for shorter commutes. Commuting times (in log) can be decomposed into the (log) distance D_{jn} from work to residential neighborhoods minus the (log) average speed S_{jmn} along the route on the chosen travel mode. The utility gain from commuting distance D_{jn} at speed S_{jmn} is denoted:

$$\alpha_{my}^S S_{jmn} - \alpha_y^D D_{jn}$$

where parameters α_{my}^S and α_y^D dictate the income group-specific preferences over speeds and distances (respectively). Note that when $\alpha_{my}^C = \alpha_y^D$, they are just the coefficient on (log) commuting time. But I allow preferences over speed α_{my}^S to also vary with the choice of travel mode m . The value of travel time spent riding the transit may differ from the time spent driving (or walking), and consequently, so may preferences for travel times savings on each travel mode (and differentially across income groups).

On the other hand, parameters α_y^D reflect the net gains from shorter commutes unconditional on mode choice. Since workers have different work locations within the

city, they differ in their distances to high quality residential neighborhoods and, consequently, in their accessibility gains from a longer commuting distance. So, α_y^D also encapsulate differences in the geography of high- and low-income jobs within the city. Workers from an income group with more jobs farther away from desirable neighborhoods are likely to be more willing to commute longer and have a smaller α_y^D .¹⁸

Neighborhoods differ in their supply of developable land and a competitive housing market determines the equilibrium housing prices p_n (per unit of space) faced by the neighborhood's residents. While prices depend on the aggregate demand for housing space in each neighborhood, each worker takes these prices as given when making housing consumption and location choices. Housing is a normal good and individual demand for housing space is increasing with income and decreasing with the price of housing. More specifically, conditional on residing in neighborhood n , the housing consumption of worker i is:

$$h(p_n, w_i) = (p_n)^{\alpha_h} (w_i)^{\alpha_w} \quad (4.1)$$

where $\alpha_h < 0$ is the price elasticity and $\alpha_w > 0$ is the income elasticity of housing demand.

Net of preferences over housing costs and commuting times, each worker's preferences over neighborhoods and travel modes can be decomposed into two components: a common preference across all workers in an income group and an idiosyncratic preference. Let δ_{mny} denote the income group-specific utility from choosing neighborhood n and travel mode m . This utility shifter captures differences across modes in the monetary cost of travel (such as of vehicle ownership or transit fare) that affects each income group differently. They also capture differences in the quality of residential amenities (such as schooling and crime) as well as in location-specific attributes of travel (such as convenience of parking or waiting at the nearest transit stop). The latter may include differences in how well (on average) the commuting mode connects the residential neighborhood to non-commuting destinations and non-residential amenities such as restaurants and shopping malls. While the (unobserved) mode choices on non-commuting trips may be different from the observed mode choice on commutes, the gains from owning a vehicle or a bus pass are greater when they improve access to

¹⁸Alternatively, if jobs are more substitutable across space (e.g. in terms of wages) for one income group, they may have a stronger preference for more centrally located jobs and a higher α_y^D . Modeling the geography of jobs (and work location choices) explicitly is beyond the scope of this paper.

more than just the immediate work location.

Workers also have idiosyncratic preferences over neighborhood-mode alternatives and I let ϵ_{imny} denote their idiosyncratic utility gains from choosing neighborhood n and mode m . Assume ϵ_{imny} are random draws from a type 1 extreme value (T1EV or Gumbel) distribution that is identical across workers and independent of their commuting and housing preferences. Together with the aforementioned deterministic components of utility, workers' choices in equilibrium maximize the following (indirect) utility function:

$$U_{mn|ijy} \equiv \alpha_{my}^S S_{jmn} - \alpha_y^D D_{jn} + \frac{(w_i)^{1-\alpha_w}}{1-\alpha_w} - \frac{(p_n)^{1+\alpha_h}}{1+\alpha_h} + \delta_{mny} + \epsilon_{imny} \quad (4.2)$$

The parameters α_w and α_h determine the diminishing marginal utility from higher incomes and the marginal disutility from higher prices (respectively). The housing demand function in (4.1) follows from Roy's Identity.¹⁹

Preference parameters α_{my}^S , α_y^D , α_w and α_h may vary across cities, but I drop the city subscripts to simplify notation. That means preferences over commuting and housing depend on city-level attributes such as (but not limited to) the spatial distribution of high- and low-income jobs with respect to the travel network and city-level housing constraints. These city-level attributes are exogenous with respect to each worker's decision problem.

Finally, given the distribution of the logit error term ϵ_{imny} , the probability of a worker from income group y and work location j choosing mode m and neighborhood n is

$$\pi_{mn|jy} = \frac{\exp(V_{mn|jy})}{\sum_{m' \in M} \sum_{n'} \exp(V_{m'n'|jy})} \quad (4.3)$$

$$\text{where } V_{mn|jy} \equiv \alpha_{my}^S S_{jmn} - \alpha_{my}^D D_{jn} + \delta_{mny} - \frac{(p_n)^{1+\alpha_h}}{1+\alpha_h}$$

4.2 Identification

In applying this utility specification to data, I address three important empirical challenges to identifying preferences for faster commutes. First, travel times on commutes

¹⁹By Roy's Identity:

$$h(p, w) = -\frac{dU/dp}{dU/dw}$$

depend on both the spatial distribution of transportation infrastructure (such as transit routes and highways) and the spatial distribution of jobs relative to residential neighborhoods. If work locations for higher income groups are farther away from desirable residential neighborhoods, they may appear to have a smaller disutility from longer commutes despite having a higher opportunity cost of travel time. Decomposing commuting times into shortest-route road distances D_{jn} and mode-specific speeds S_{jmn} allows me to isolate the two effects and identify preferences over access to faster commutes conditional on proximity to jobs. Furthermore, conditional on the income group-specific fixed effects, I am identifying the coefficient on speed using variation across individuals in the same income group.

Second, commuting speeds may be correlated with other (unobservable) attributes of residential neighborhoods and travel modes. For example, if transit planners are more likely to expand high-speed transit routes into neighborhoods with attributes more desirable to the rich, then unless I control for these correlated neighborhood attributes in my regression, higher income commuters would appear to have a higher coefficient on transit speed than they actually do. My inclusion of alternative-specific constants for each income group δ_{mny} (fixed effects) essentially controls for preferences over unobservable neighborhood-mode attributes.

Third, commuting speeds may be systematically correlated with the locations of high- and low-income jobs. For example, if work locations of some income groups are better connected by high-speed transit than driving relative to the work locations of others, then these income groups would appear to have a higher coefficient on transit speed than they actually do. To address this concern, I standardize the commuting speeds and distances faced by each worker to mean 0 (and standard deviation 1) conditional on travel mode. In doing so, any mean preference for one travel mode over another within an income group y is absorbed by the group's corresponding alternative-specific constant δ_{mny} . So, conditional on commuting distance and income group-alternative-specific constants, the coefficients on speed α_{my}^S are the gains from shorter commuting time identified off of mode-specific variation in speeds to different work locations in the city.²⁰

In addition to the coefficients on commuting speed, I need to estimate housing demand parameters α_w and α_h to be able to compare preferences for access to faster commutes in terms of workers' willingness to pay for housing. This exercise poses two

²⁰Later on, I transform the speeds and distances back to their unstandardized levels for evaluating willingness to pay and transit ridership responses to a percent change in travel speeds.

additional econometric challenges. First, the price elasticity of housing demand α_h is not identifiable from the utility specification because housing prices p_n are necessarily correlated with neighborhood characteristics captured by the fixed effects δ_{mny} . Second, the choice probabilities in (4.3) do not inform us at all about the income elasticity of housing demand α_w . So, I need to identify these parameters separately. To do so, I exploit the housing expenditure patterns of a representative micro-sample of each city’s working population. Consider the log of the housing demand function in (4.1), which I can rewrite as a linear relationship between the log of total housing expenditure as a share of income (on the left) and the logs of housing prices and incomes (on the right):

$$\ln\left(\frac{h_{in}p_n}{w_i}\right) = (1 + \alpha_h)\ln(p_n) + (\alpha_w - 1)\ln(w_i) \quad (4.4)$$

Then the price and income elasticities of housing demand follow directly from the coefficients of (log) price and (log) income above.

4.3 Estimation

Estimation of the model parameters proceeds separately for each city and in two stages. The first stage estimates the housing demand parameters α_w and α_h using micro-data on individual housing expenditures in an OLS estimation based on (4.4). For each worker in the census micro-sample, I observe both precise household incomes and the share of that income spent on housing expenditures. I can combine them with tract-level standardized housing prices from Davis et al. [2020].²¹ Then I regress the log of housing expenditure share on log housing price and log household income as below.

$$\ln\left(\text{HousingExpShare}\right) = \bar{\alpha}_h \ln\left(\text{Price}\right) + \bar{\alpha}_w \ln\left(\text{Income}\right) \quad (4.5)$$

where the coefficients are $\bar{\alpha}_h = 1 + \alpha_h$ and $\bar{\alpha}_w = \alpha_w - 1$. See Appendix Section A.3 for estimation details and results. Having estimated α_h , the housing price component of each worker’s choice probabilities $-(p_n^{1+\alpha_h})/(1 + \alpha_h)$ is just a neighborhood-specific constant from here on.

The second stage estimates parameters α_{my}^S and α_y^D together with fixed effects δ_{mny}

²¹I do not observe the workers’ tracts of residence in the microdata. The smallest known geography of residence is the PUMA, which are slightly larger. So, instead, I assign each worker the expected housing price experienced by workers in the same income bin and PUMA. See Appendix Section A.3 for details.

using the data on observed commuting flows and counterfactual travel times in a maximum likelihood estimation based on (4.2). The estimation maximizes the probability that the model correctly matches each worker in the city to their observed neighborhood and mode in the CTPP data. In particular, estimated parameters maximize the following sum across all workers of the log-likelihood of their observed choices:

$$L = \sum_y \sum_j \sum_n \sum_{m \in M} P_{jmn_y} \ln(\pi_{mn|jy}) \quad (4.6)$$

where P_{jmn_y} is the observed population of commuters in income group y and work location j who choose mode m and residence n . The estimation procedure then consists of numerically searching over the twelve α_{my}^S parameters and the four α_y^D parameters as well as the full matrix of fixed effects δ_{mny} in order to maximize L .

The set of work locations are the census tracts in the city that receive non-zero commutes. The choice set of residential neighborhoods in each city is the set of census tracts with non-zero observed population of workers.²² The number of residential tracts ranges from 58 in my smallest city (Trenton, NJ) to 3050 in my largest (New York, NY). So, given the large number of fixed effects to be estimated for every mode, neighborhood and income group combination, I exploit a contraction mapping approach popularized by Berry et al. [1995] to speed up convergence to the optimal parameter estimates. See Appendix Section A.4 for details.

Across all cities, income groups and travel modes, I estimate 588 different coefficients on commuting speed. To make them comparable across cities and income groups, I combine the estimated coefficients with my parameter estimates from the first stage to characterize preferences in terms of the implied marginal willingness to pay (MWTP) in annual housing costs for faster commutes.²³ Appendix Table A.2 reports the distribution of the (raw) estimated coefficients on commuting distance and speed (α_y^D and α_{my}^S) across the 49 cities. The following section explores how the implied MWTP varies across cities and income groups.

²²Commuters with either residence or work location outside of the extent of the city are dropped from the sample.

²³The marginal willingness to pay (MWTP) for higher commuting speed is $-\frac{dU/dS_{jmn}}{dU/dp_n}$.

5 Estimated Preferences for Faster Transit

This section presents the estimated preferences for faster transit commutes in three stages. First, I characterize the value of travel time conditional on travel mode choice. In other words, how much are transit riders willing to pay for faster commutes (compared to drivers)? Second, I characterize the marginal propensity of consumers to ride transit in response to shorter transit travel times. In other words, how do increases in transit speed affect transit ridership? Third, I combine the two results to characterize the overall expected welfare gains from increases in transit speed (unconditional on mode choices) and how these gains compare for rich and poor commuters.

5.1 Willingness to pay for faster commutes

Conditional on travel mode choices, the mean estimated MWTP (per year) for a one percent increase in travel speed on commutes (across all cities) is \$98 among transit riders and \$142 among drivers. Assuming workers commute 5 days a week and commutes make up 35% of their total time spent traveling (based on reported travel times in the 2017 NHTS)²⁴, the mean MWTP estimates for speed imply a mean value of travel time savings (VTTS) among transit riders of \$7.4 per hour (and roughly 40% of the median transit rider’s wage). In comparison, the mean VTTS among drivers is \$15.5 per hour (which is 86% of the median driver’s wage).²⁵ My mean driving estimates are similar to contemporary value-of-time estimates from other papers using alternative methodologies (Small, 2012), such as means of \$13-\$14 per hour in Prague (Bucholz et al., 2020) and Vancouver (Craig, 2019). There are no comparable estimates in the literature of the value of travel time on transit.

As shown in Table 3, the mean estimates mask large variation across cities. The table reports the mean MWTP for faster travel by mode choice and city ranked by the MWTP among transit riders.²⁶ Because of the large number of commuting trips informing these preference estimations, asymptotic standard errors are tiny (typically

²⁴The share of total travel time spent on commutes to work is calculated from the share of reported travel times spent on trips to work in the 2017 US National Household Travel Survey (NHTS). I assume increases in travel speed on commutes also increases travel speeds on all other trips at the same rate.

²⁵One reason for the VTTS among transit riders being a smaller share of wages than the VTTS among drivers is that transit riders are primarily concentrated in higher-income cities. So, across all cities, the average transit rider has a higher income than the average driver.

²⁶A table of results for the full list of cities is in the Appendix.

around one cent or less in MWTP) and omitted from the tables.²⁷ Focusing first on transit users, in San Francisco (the top ranked city on the list), the MWTP for faster transit is \$374 per year, almost four times the average across all commuters. To benchmark these magnitudes, consider MWTP estimates for other locational attributes. For example, Bayer and McMillan [2012] estimate MWTP (per year) in the San Francisco Bay area of: \$236 for access to schools with (1 standard deviation) higher average test scores, \$126 for 10% more college-educated neighborhoods and \$436 for neighborhoods with \$10,000 higher average incomes. While city-level MWTP for faster commutes for drivers are similar in magnitude to those for transit riders, the rank ordering is different. When interpreting these preference estimates, bear in mind that they reflect how aggressively transit riders and drivers bid for access to (and sort into) locations with faster commutes. So, some of these cross-city differences in mean MWTP also stem from differences in housing market constraints and urban amenities that make housing in some cities more expensive than in others.

Table 3: **Cities ranked by mean MWTP for faster transit commutes**

Rank	City	MWTP for faster transit	MWTP for faster driving
1	San Francisco, CA	\$ 374	\$ 302
2	Seattle, WA	\$ 188	\$ 179
3	New York, NY	\$ 178	\$ 345
4	San Jose, CA	\$ 169	\$ 139
5	Boston, MA	\$ 148	\$ 189
6	Washington, DC	\$ 129	\$ 156
7	Vallejo-Fairfield, CA	\$ 119	\$ 69
8	Chicago, IL	\$ 116	\$ 179
9	Los Angeles, CA	\$ 114	\$ 102
-	-	-	-
20	Miami, FL	\$ 64	\$ 75
-	-	-	-
29	Phoenix, AZ	\$ 44	\$ 47
-	-	-	-
36	Urban Honolulu, HI	\$ 32	\$ 19
-	-	-	-
49	Las Vegas, NV	\$ 9	\$ 17

Note: Cities are ranked by the mean MWTP for faster transit. MWTP values are means across all commuters for 1% change in travel speed on their observed commutes (i.e. conditional on commuters choosing their observed modes and neighborhoods). See mean MWTP estimates for full list of cities in the Appendix.

²⁷In work in progress, I bootstrap the standard errors with Monte Carlo simulations to derive more credible estimates.

Table 4: Mean MWTP for 1% increase in commuting speed

City	Mode	< \$35k	\$35k-\$50k	\$50k-\$75k	>\$75k
New York	transit	\$43	\$69	\$125	\$219
	driving	\$134	\$211	\$273	\$401
Los Angeles	transit	\$42	\$62	\$80	\$146
	driving	\$51	\$81	\$81	\$123

Note: MWTP values are means across all commuters for 1% increase in travel speed on their observed commutes (i.e. conditional on commuters choosing their observed modes and neighborhoods). Asymptotic standard errors are less than a cent.

Table 5: Mean relative MWTP across all cities

Mode	< \$35k	\$35k-\$50k	\$50k-\$75k	>\$75k
transit	1.00	1.40	2.00	3.24
driving	1.00	1.47	1.86	2.57


Note: Reported values are the MWTP estimates for 1% increase in commuting speed divided by the lowest income group's MWTP and averaged over commutes across all cities.

Incomes are an important determinant of a commuter's MWTP estimate. Table 4 decomposes the mean MWTP by commuter's income bracket and reports it for the two largest cities in my sample. Unsurprisingly, richer transit riders have higher MWTP for faster transit commutes than poorer transit riders. Also, richer drivers have higher MWTP for faster driving commutes than poorer drivers. This pattern is consistent with the rich having a higher overall value of travel time savings than lower income commuters. When I aggregate my estimates across all commuters, the magnitude of the income differences are comparable to extant reduced form estimates in the literature.²⁸ Some of the income elasticity of MWTP are undoubtedly due to differences in the mean *ability* to pay (and richer commuters generally spending more on housing). More notably, based on the differences across income brackets, the income elasticity of the demand for faster travel appears to be higher among transit riders than drivers. Table 5 pools together commuters across all 49 cities and, for comparability across cities, presents the MWTP estimates as multiples of the lowest income group's MWTP. The average transit commuter with income over \$75,000 is willing to pay over three times more for a one percent increase in commuting speed than the average transit commuter with income below \$35,000. This is not just driven by differences in the ability to pay. The table also shows the relative MWTP across income groups among drivers, and the

²⁸Small [2012] reviews contemporary empirical estimates of value of time savings (VTTS) on commutes and cites income elasticities of VTTS typically between 0.5 and 0.7.

income elasticity of the willingness to pay for faster commutes is much smaller among drivers than among transit riders (like in New York and Los Angeles).

Figure 5.1a plots these ‘relative’ MWTP estimates by city against mean transit speeds relative to driving (from column 4 of Table 2). The rich have higher MWTP for increase in transit speed relative to the poor in cities with faster transit. This figure highlights a key finding of my work: transit improvements are relatively more attractive to the rich when transit is fast. Another dimension of transit which is often associated with use by the wealthy is rail transit versus bus transit. Rail transit typically has higher velocity than buses, so the rail composition of a city’s transit network can proxy for average travel velocity on transit (and an alternative to my measure of travel speeds). Figure 5.1b plots the ‘relative’ MWTP estimates against each city’s rail share of transit usage (from column 6 of Table 2). Transit improvements are relatively more attractive to the rich when transit is more rail-based.



(a) by mean transit speeds (relative to driving) (b) by rail share of transit commutes in the city

Figure 5.1: Mean MWTP for 1% increase in transit speed (relative to lowest income group). Each observation corresponds to a city. Vertical axis depicts the MWTP for faster transit as a fraction of the MWTP of commuters with incomes less than \$35,000 (indicated by solid black line at 1). Horizontal axis depicts either (a) the ratio of driving to transit travel times (across all observed commutes) in the city or (b) the share of transit riders in the city who commute by rail transit. Confidence intervals for each linear fit are shaded in corresponding color. For commuters with incomes \$35k-\$75k, figures plot population-weighted means of the MWTP estimates for the two middle-income groups in my data.

5.2 Willingness to ride transit

So far, I have focused on heterogeneity in MWTP for transit speed (or, if you will, the demand for faster travel) among commuters who ride transit. MWTP is central to evaluating welfare from the perspective of transit users, but transit policy is often predicated on a broader set of concerns including reducing congestion and climate change concerns. Evaluating policy along these dimensions requires assessing how policies impact the decision to use transit instead of driving. To this end, I use my model to estimate the probability of non-transit commuters switching to riding transit if transit is made faster. Let R_{jny} denote the transit ridership among income group y on commutes between neighborhood n and work location j . I can solve for a commuter’s marginal willingness to ride transit (or, in aggregate terms, the marginal change in transit ridership), denoted MWTT, in response to a percent increase in transit speed on their commuting route.²⁹ The MWTT measures the predicted change in the probability of transit use along a given work-residence commuting route in response to a percent increase in speed along the route.

Table 6 reports the mean MWTT across all commuters in a city conditional on observed residential location choices. The cities at the top of the list, where marginal improvements in transit speed would be most effective at generating new transit ridership, are likely to be cities with high pre-existing transit ridership (but not always). The top of the list includes both cities with high rail transit usage among transit riders (such as Chicago, Washington and Boston) and ones with very low rail transit usage (such as Seattle, Portland and Pittsburgh).

These cities also attract transit riders at different rates across income groups. Table 7 compares the MWTT across income groups for New York and for Los Angeles. In New York, a one percent increase in transit speeds everywhere increases transit ridership more among higher income commuters than lower income commuters. However, the opposite is true in Los Angeles, where lower income commuters are twice as likely to increase transit ridership. The case of Los Angeles is more common among other cities, but there is also a generalizable pattern here. Cities with high (baseline) transit ridership among commuters are more likely to have high MWTT among richer

²⁹Formally, the marginal willingness to ride transit (MWTT) is defined:

$$MWTT_{jny} = \frac{dR_{jny}}{dS_{jtransitn}} \equiv \frac{d}{dS_{jtransitn}} \left(\frac{\pi_{transitn|jy}}{\sum_{m \in M} \pi_{mn|jy}} \right) = \alpha_{transity}^S R_{jny} (1 - R_{jny})$$

Table 6: **Cities ranked by MWTT from 1% increase in transit speeds**

Rank	City	%age pt change in transit ridership	Baseline transit ridership (in %)
1	San Francisco, CA	3.52	15.5
2	Chicago, IL	2.97	11.9
3	Washington, DC	2.97	14.5
4	Seattle, WA	2.88	8.6
5	Boston, MA	2.76	12.4
6	Portland, OR	2.22	6.6
7	New York, NY	2.22	30.7
8	Philadelphia, PA	1.81	9.7
9	Pittsburgh, PA	1.77	6.0
-	-	-	-
18	Miami, FL	0.80	3.8
-	-	-	-
25	Urban Honolulu, HI	0.71	8.2
-	-	-	-
34	Phoenix, AZ	0.55	2.3
-	-	-	-
49	Rochester, NY	0.13	2.1

Note: Cities are ranked by the MWTT in response to a 1% increase in transit speed along all observed commutes (i.e. conditional on commuters choosing their observed neighborhoods). See mean MWTT for full list of cities in the Appendix.

commuters (relative to the MWTT among poorer commuters). Table 8 groups together cities by each city’s (baseline) transit ridership. For comparability across cities, I present the MWTT estimates as fractions of the lowest income group’s MWTT. In most cities, baseline transit ridership is low and poorer commuters have much larger MWTT. However, in the five cities where more than 10% of the commutes are by transit, the MWTT is similar if not larger for richer commuters.

These five cities (New York, San Francisco, Chicago, Boston and Washington DC) also happen to be (among the seven) cities with more rail transit riders than road transit riders. Since higher income transit riders benefit more from improvements in rail-heavy transit networks (as I showed in Section 5.1), it is unsurprising that these improvements also increase transit usage relatively more among higher income commuters.

Table 7: **Mean MWTT from 1% increase in commuting speed**

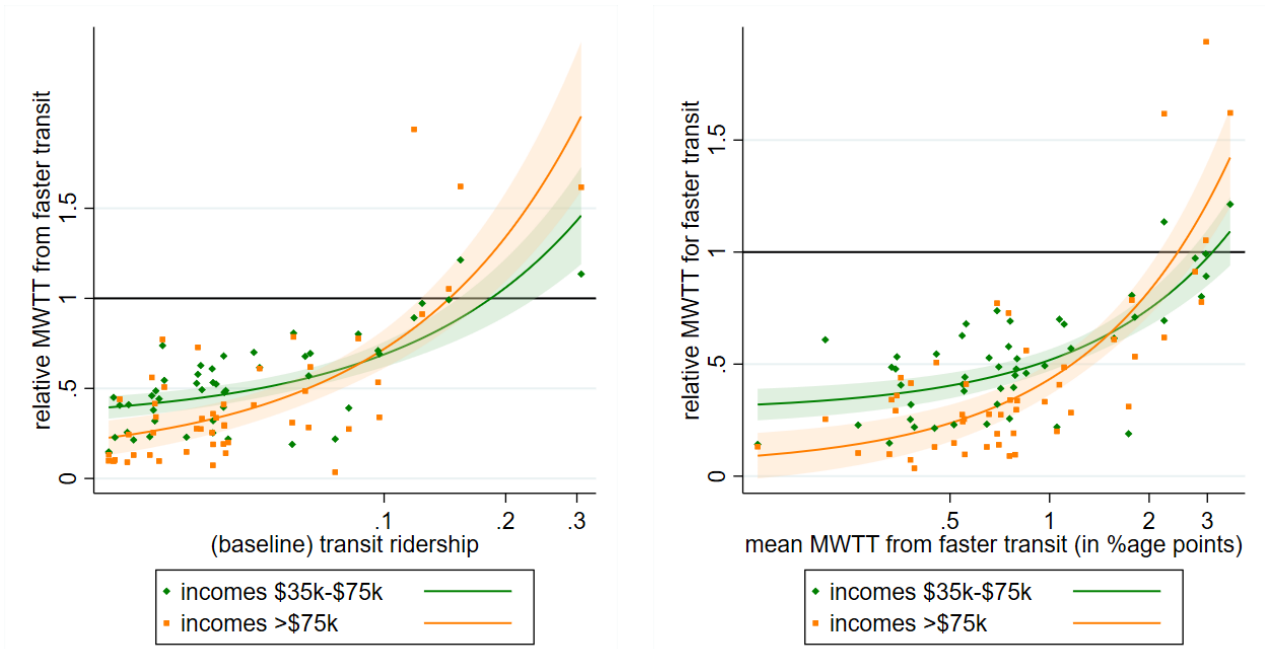
City	< \$35k	\$35k-\$50k	\$50k-\$75k	>\$75k
New York	1.5%	1.3%	1.9%	2.5%
Los Angeles	1.8%	1.4%	1.2%	0.9%

Note: Table reports percentage point change in total transit ridership across all commuters in response to 1% increase in transit speeds.

Table 8: Mean relative MWTT across cities

Cities with . . .	< \$35k	\$35k-\$50k	\$50k-\$75k	>\$75k
less than 10% commuters riding transit	1.00	0.61	0.52	0.40
more than 10% commuters riding transit	1.00	0.86	1.15	1.53

Note: Reported values are means across all commuters of their MWTT estimate divided by the lowest income group's MWTT over the same commuting route.



(a) by (baseline) share of commuters riding transit

(b) by overall MWTT (across all commuters)

Figure 5.2: **Mean MWTT from 1% increase in transit speed (relative to lowest income group)**. Each observation corresponds to a city. Vertical axis depicts the MWTT for faster transit as a fraction of the MWTT of commuters with incomes less than \$35,000 (indicated by solid black line at 1).

Horizontal axis depicts in log scale either (a) the baseline transit ridership (across all observed commutes) in the city or (b) the mean MWTT across all commuters. Confidence intervals for each linear fit are shaded in corresponding color. For commuters with incomes \$35k-\$75k, figures plot population-weighted means of the MWTT estimates for the two middle-income groups in my data.

More generally, cities with higher overall (baseline) transit ridership and higher overall marginal transit ridership (namely higher overall MWTT) have larger relative MWTT among higher income commuters. Figure 5.2 illustrates this point by plotting the relative MWTT across income groups as a function of both the city's overall baseline

transit ridership and overall MWTT.³⁰ In cities with higher baseline transit ridership and where transit improvements are most effective at generating new ridership, transit improvements also increase ridership relatively more among higher income commuters.

Within cities

The data show similar patterns in marginal transit ridership across locations *within* cities. So far, the analysis has focused on city-wide improvements in transit speeds. For evaluating MWTT within cities, I now consider increasing speeds only along particular commuting routes. To illustrate general patterns, I plot means for the route-specific results aggregated along two dimensions: (standardized) driving speed and (standardized) transit speed.³¹ I plot these results for commuters in each of three income brackets. Figure 5.3 shows contour plots for New York and Los Angeles of the MWTT from an increase in transit speed at different points along the city’s observed commuting network.³² The x- and y-axes depict existing driving and transit speeds on the route. The axes scales are fixed so that the colors representing the MWTT are comparable across income groups and cities.


I highlight two regularities that are clear from these graphs. First, as seen from the increasingly reddish shades at the top-left of each graph, the marginal gains in transit ridership are higher along routes where transit is already relatively fast (or driving is relatively slow). The figures suggest that the relationship between transit ridership and transit speed is convex. Marginal transit improvements may seem ineffective at the beginning when transit is slow, but would yield increasingly larger ridership returns.

Second, in New York, the ridership gains among higher income commuters are much larger (compared to lower income commuters) where driving is relatively slow. Whereas in Los Angeles, it is the opposite: lower income commuters are the ones more likely to increase transit ridership along the (relatively) slow driving routes. The graphs for other rail-transit cities like Chicago, Washington DC and San Francisco with high overall transit ridership resemble that of New York in that high-income commuters have


³⁰Income differences in mean MWTT across cities are also correlated with mean transit speeds relative to driving and the city’s rail share of transit usage. Appendix Figures A.3 plot the relative MWTT across income groups as a function of both the city’s rail share of transit commutes and its mean transit speed relative to driving.

³¹As before, for comparability across cities, speeds are normalized across all commutes within each city.


³²Note that the MWTT does not capture commuters moving across residential neighborhoods, and hence the graphs depict the change in transit ridership among commuters given their (observed) work and residence. Section 5.3 relaxes this assumption.




(a) New York, Incomes <\$35k




(b) New York, Incomes \$35k-\$75k




(c) New York, Incomes >\$75k



(d) Los Angeles, Incomes <\$35k



(e) Los Angeles, Incomes \$35k-\$75k



(f) Los Angeles, Incomes >\$75k

Figure 5.3: Mean MWTT by location of transit improvement. The x- and y-axes depict standardized driving and transit speeds (resp.) on the commuting route. The z-axis depicts the percentage point change in transit ridership in response to a 1% increase in transit speed along the commuting route. The z-axis colours are fixed across all graphs. Speeds are standardized (to mean 0 and std. dev. 1) across trips between observed work-residence pairs within each CBSA... Trips at the top and bottom percentiles of speeds are ignored. White spaces in the graphs correspond to 0.1-by-0.1 cells with fewer than 20 commutes.

a larger MWTT where driving is relatively slow. Whereas, graphs for most other cities look like those of Los Angeles where low-income commuters have a larger MWTT when driving is relatively slow.³³ There are, however, exceptions like Seattle (where transit is fast and only 5% of transit commutes are by rail transit) where the distribution of the ridership gains for high and low income commuters look similar (see Appendix Figure A.4).

More generally, the routes where marginal transit improvements are most effective at generating new transit ridership across all income groups (such as where driving is relatively slow to begin with) are ones: (a) where the ridership gains are larger among the rich in high-speed rail-transit cities like New York or (b) where the gains are larger among the poor in low-speed road-transit cities like Los Angeles.

5.3 Distribution of welfare gains

Finally, what are the welfare gains across income groups from faster transit commutes? In Section 5.1, I presented the gains from faster travel for transit riders and drivers conditional on their observed mode and location choices. Now, having characterized how higher transit speeds affect the probability of riding transit (conditional on neighborhood choices), this section quantifies the average commuter’s expected marginal gains (in terms of their marginal willingness to pay) from increase in travel speeds *unconditional* on their mode and location choice, denoted unconditional MWTP or ‘uMWTP’ for short. In other words, I compute the marginal gains from increases in transit speed for all commuters (not just transit riders) accounting for re-sorting across both travel modes and residential locations.³⁴

Table 9 compares estimates of uMWTP for a one percent increase in travel speeds by driving and transit in New York and Los Angeles. The uMWTP for an increase in transit speeds are an order (or two) of magnitude smaller than the uMWTP for an increase in driving speeds, which is unsurprising given generally low baseline transit ridership. Higher income commuters have a higher willingness to pay than lower income commuters for faster driving commutes, but the income elasticity of the gains from

³³While the increase in transit ridership (in terms of percentage point change) is larger among lower-income commuters, the percentage change from baseline transit ridership is larger among higher income commuters, who have very low transit usage in cities like Los Angeles.

³⁴Note that welfare gains in this context only refer to the direct utility gains from shorter commuting times as formalized in Section 4. They do not account for general equilibrium effects, such as through changes in congestion or the locations of jobs and residential amenities.

faster transit commutes varies by city. In New York, the uMWTP for faster transit is higher among richer commuters. Whereas in Los Angeles, the uMWTP for faster transit is higher among poorer commuters. Table 10 generalizes this result across all cities and presents the uMWTP estimates relative to that of the lowest income group’s. As in Los Angeles and New York, higher income commuters consistently benefit more from increases in driving speeds. However, the lowest income commuters benefit most on average from increases in transit speeds. And as with transit ridership in the previous section, the gains from increases in transit speeds are (over four times) larger for poorer commuters in cities with low baseline transit ridership but (over two times) larger for richer commuters in cities with high baseline transit ridership.

Table 9: Mean u(nconditional)MWTP for 1% increase in travel speeds

City	Mode	< \$35k	\$35k-\$50k	\$50k-\$75k	>\$75k
New York	transit	\$18	\$24	\$33	\$44
	driving	\$131	\$231	\$316	\$478
Los Angeles	transit	\$2.9	\$2.9	\$2.2	\$0.9
	driving	\$62	\$101	\$101	\$154

Note: uMWTP values are means across all commuters in the income group for 1% change in travel speeds everywhere. Asymptotic standard errors are less than a cent.

Table 10: Mean relative uMWTP across all cities

	Mode	< \$35k	\$35k-\$50k	\$50k-\$75k	>\$75k
All cities	driving	1.00	1.56	2.00	2.83
	transit	1.00	0.84	0.83	0.92
with less than 10% transit ridership	transit	1.00	0.70	0.48	0.24
with more than 10% transit ridership	transit	1.00	1.13	1.54	2.33

Note: uMWTP estimates are divided by the lowest income group’s and averaged over commutes across all cities.

How do these distributional effects compare to the overall welfare gains from faster transit? Table 11 ranks cities by their overall mean uMWTP for faster transit (across all income groups).³⁵ Cities with higher overall uMWTP (across all commuters) for faster transit are ones where both overall (baseline) transit ridership and the rail share

³⁵While some of the cross-city differences in the magnitudes of uMWTP may be attributable to city-specific housing markets, cities with higher uMWTP for faster transit also have higher uMWTP for transit relative to driving. Column 6 of the table presents the ratio of the uMWTP for faster transit to the uMWTP for faster driving, and a ranking of cities based on this ratio is strongly correlated to the ranking presented. See Appendix for a complete ranking of cities by uMWTP for faster transit.


of transit ridership are high. Most remarkably, the five cities with the highest uMWTP are also the only cities in my sample where commuters with incomes above \$75,000 have a larger uMWTP for faster transit than commuters with incomes below \$35,000. These cities are able to attract disproportionately more high income transit riders (as seen in Section 5.2), and higher income transit riders have a higher willingness to pay for faster commutes especially when transit is already relatively fast and rail transit is more prevalent (as seen earlier in Section 5.1).

More generally, this result reflects the fact that as I move up the ranking of cities, richer commuters stand to benefit increasingly more (relative to poorer commuters) from marginal improvements in transit speed. Figure 5.4a illustrates this point by plotting each city's mean uMWTP for faster transit of higher income commuters relative to commuters with incomes below \$35,000. The horizontal axis depicts the mean uMWTP across all commuters (in log scale). Cities with the highest per capita gains from marginal transit improvements are also ones where the welfare gains are more likely to accrue to the rich. And these are also the cities where transit improvements are most effective at generating new transit ridership, as shown in Figure 5.4b. After all, overall gains in transit ridership are higher when, as shown earlier in Section 5.2, the gains are also disproportionately higher for the rich than the poor.

Table 11: Cities ranked by uMWTP for faster transit

Rank	City	uMWTP for faster transit			Relative to driving	
		all commuters	incomes <\$35k	incomes >\$75k	all commuters	commuters
1	San Francisco, CA	\$ 39.82	\$ 17.24	\$ 46.26	0.118	
2	New York, NY	\$ 38.75	\$ 18.22	\$ 44.21	0.095	
3	Boston, MA	\$ 16.19	\$ 12.90	\$ 16.04	0.075	
4	Washington, DC	\$ 14.81	\$ 13.86	\$ 14.31	0.087	
5	Chicago, IL	\$ 11.44	\$ 4.81	\$ 16.29	0.053	
6	Seattle, WA	\$ 7.87	\$ 9.33	\$ 6.96	0.038	
7	Philadelphia, PA	\$ 5.69	\$ 8.20	\$ 4.49	0.038	
8	Portland, OR	\$ 2.49	\$ 4.26	\$ 1.67	0.023	
9	Pittsburgh, PA	\$ 2.01	\$ 3.51	\$ 1.52	0.018	
-	-					
21	San Diego, CA	\$ 0.51	\$ 1.37	\$ 0.18	0.005	
-	-					
35	Phoenix, AZ	\$ 0.24	\$ 0.65	\$ 0.10	0.004	
-	-					
49	Provo-Orem, UT	\$ 0.02	\$ 0.06	\$ 0.00	0.001	

Note: Cities are ranked by their mean uMWTP (across all commuters) for 1% increase in transit speeds. Reported uMWTP values are estimates of mean MWTP across all commuters unconditional on their choices of mode and neighborhood. Ratio in column 6 divides uMWTP estimates in column 3 by estimates of the city's mean uMWTP for 1% increase in driving speeds. See ranking for full list of cities in the Appendix.



(a) by mean uMWTP across all commuters

(b) by mean MWTT across all commuters

Figure 5.4: Mean unconditional MWTP for a 1% increase in transit speed.

Vertical axis depicts the income group's mean (relative to commuters with income below \$35,000) and horizontal axis depicts (in log scale) the mean across all commuters of either (a) uMWTP or (b) MWTT. Each observation corresponds to a city. Confidence intervals for each linear fit are shaded in corresponding color. For incomes \$35k-\$75k, I plot population-weighted averages of the uMWTP of the two middle-income groups in my data.

6 Conclusion

In this paper, I introduce a methodology for evaluating the demand for faster commutes by public transit and driving based on observed residential location and travel mode choices within cities. In doing so, I address two important empirical challenges that has limited past work on this topic. The first one is a (sparse) data challenge: I need to compare chosen (and observed) commutes to unchosen (and unobserved) ones. To measure the latter, I combine millions of scraped trip queries on Google Maps with data on street networks to predict travel times on all possible alternative commutes between census tracts in US cities. The second challenge is to disentangle the extent to which observed choices and the gains from them (as reflected in housing prices) are due to differences in commuting speeds as opposed to other spatially correlated features of travel modes and neighborhoods. To that end, I propose a discrete choice model that complements my rich data environment with detailed fixed effects in order to identify heterogeneous preferences over commuting speeds. Applying this model to 49 US cities with different transit networks reveals many new insights on the expected ridership and welfare gains from transit improvements across income groups and cities.

Among other things, I show that the demand for faster transit commutes is small relative to the demand for faster driving commutes and depends importantly on the speed of transit relative to driving along commutes as well as on the prevalence of rail transit in the city. Ridership and welfare gains from transit improvements are larger for high income commuters in cities with already high transit ridership, relatively fast transit and high rail transit usage. The opposite is true (that is, larger gains for lower income commuters) in cities with low baseline transit ridership, relatively slow transit and low rail transit usage. And because higher income transit riders have a higher willingness to pay for faster transit commutes, cities where transit improvements are more attractive to the rich are also the ones where they generate more overall transit ridership and welfare. While transit improvements are often believed to reduce inequality in cities, this result suggests that transit improvements most in demand (and, consequently, more likely to be cost-effective and to be realized) are likely to trade off equity for efficiency.

While the paper’s findings shed light on several important policy questions, it also opens up new ones that the paper leaves unanswered. For instance, why are transit improvements in rail-transit-heavy cities more likely to benefit the rich? One hypothesis is that because rail transit expansions can be much costlier than road transit expansions,

transit planners may be under greater pressure to target efficiency (and high income commuters) over equity (and low income commuters) when choosing where to improve rail transit. Whereas, with buses, planners may focus more on equity. They may also be wary of transit-induced neighborhood gentrification and income segregation in the city. In ongoing work in progress, I am simulating the effect of counterfactual transit improvements in US cities on residential location choices in order to study who are likely to gentrify newly transit-accessible neighborhoods.

References

- Prottoy A. Akbar. Public transit access and income segregation. Working Paper, University of Pittsburgh, 2020.
- Prottoy A. Akbar, Victor Couture, Gilles Duranton, and Adam Storeygard. Mobility and congestion in urban india. October 2018.
- Michael L. Anderson. Subways, strikes, and slowdowns: The impact of public transit on traffic congestion. *American Economic Review*, 104(9):2763–96, September 2014.
- Nathaniel Baum-Snow and Matthew E. Kahn. The effects of new public projects to expand urban rail transit. *Journal of Public Economics*, 77:241–263, 2000.
- Nathaniel Baum-Snow and Matthew E. Kahn. Effects of urban rail transit expansions: Evidence from sixteen cities, 1970-2000. Technical report, Brookings-Wharton Papers on Urban Affairs, 2005.
- Patrick Bayer and Robert McMillan. Tiebout sorting and neighborhood stratification. *Journal of Public Economics*, 96:1129–1143, 2012.
- Patrick Bayer, Robert McMillan, and Kim Reuben. What drives racial segregation? new evidence using census microdata. *Journal of Urban Economics*, 56(3):514–35, 2004.
- Patrick Bayer, Fernando Ferreira, and Robert McMillan. A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638, August 2007.

- Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, July 1995.
- David R. Bowes and Keith R. Ihlanfeldt. Identifying the impacts of rail transit stations on residential property values. *Journal of Urban Economics*, 2001.
- Nicholas Bucholz, Laura Doval, Jakub Kastl, Filip Matejka, and Tobias Salz. The value of time: Evidence from auctioned cab rides. August 2020.
- Robert Cervero and Chang Deok Kang. Bus rapid transit impacts on land uses and land values in seoul, korea. *Transport Policy*, 18(1):102–116, 2011.
- Andrea Craig. Commute mode and location choice. Working Paper, University of Windsor, 2019.
- Morris A. Davis, William D. Larson, Stephen D. Oliner, and Jessica Shui. The price of residential land for counties, zip codes, and census tracts in the united states. Federal Housing Finance Agency Staff Working Paper 19-01, February 2020.
- Arya Gaduh, Tadeja Gracner, and Alex Rothenberg. Improving mobility in developing country cities: Evaluating bus rapid transit and other policies in jakarta. Working Paper, 2020.
- Nicolas Gendron-Carrier, Marco Gonzalez-Navarro, Stefano Polloni, and Matthew A. Turner. Subways and urban air pollution. October 2020.
- Edward L. Glaeser, Matthew E. Kahn, and Jordan Rappaport. Why do the poor live in cities? the role of public transportation. *Journal of Urban Economics*, 63(1):1–24, 2008.
- Ariel Goldszmidt, John A. List, Robert D. Metcalfe, Ian Muir, V. Kerry Smith, and Jenny Wang. The value of time in the united states: Estimates from nationwide natural field experiments. 2020.
- Marco Gonzalez-Navarro and Matthew A. Turner. Subways and urban growth: Evidence from earth. *Journal of Urban Economics*, 108:85–106, 2018.
- Yizhen Gu, Chang Jiang, Junfu Zhang, and Ben Zou. Subways and road congestion. *American Economic Journal: Applied Economics (forthcoming)*, 2019.

- Arpit Gupta, Sijn Van Nieuwerburgh, and Constantine Kontokosta. Take the q train: Value capture of public infrastructure projects. Working paper, 2020.
- Eric Hannon, Jan Tijs Nijssen, Sebastian Stern, and Ben Sumers. Transit investments in an age of uncertainty. Technical report, McKinsey & Company, March 2020.
- Stephan Hebllich, Stephen J. Redding, and Daniel M. Sturm. The making of the modern metropolis: Evidence from london. *Quarterly Journal of Economics*, 135(4): 2059–2133, 2020.
- Kilian Heilmann. Transit access and neighborhood segregation: Evidence from the dallas light rail system. *Regional Science and Urban Economics*, 73:237–250, 2018.
- Matthew E. Kahn. Gentrification trends in new transit-oriented communities: Evidence from 14 cities that expanded and built rail transit systems. *Real Estate Economics*, 2007.
- Ed Kalachek. Ghetto dwellers transportation and employment. In *Conference on Transportation and Poverty*, Brookline, MA, 1968.
- Florence Kondylis, Arianna Legovini, Kate Vyborny, Astrid Zwager, and Luiza Andrade. Demand for safe spaces: Avoiding harassment and stigma. World Bank working paper (WPS 9269), 2020.
- Eunjee Kwon. Why do improvements in transportation infrastructure reduce the gender gap in south korea? Working paper, 2020.
- Yatang Lin. Travel costs and urban specialization patterns: Evidence from china’s high speed railway system. *Journal of Urban Economics*, 98(C):98–123, 2017.
- Daniel McFadden. *Spatial Interaction Theory and Planning Models*, chapter Modelling the Choice of Residential Location, pages 75–96. North Holland, 1978.
- Ismir Mulalic, Ninette Pilegaard, and Jan Rouwendal. Does improving public transport decrease car ownership? evidence from the copenhagen metropolitan area. *Regional Science and Urban Economics*, 83:1–19, July 2020.
- OECD. Infrastructure investment (indicator), 2019.

- Rahul Pathak, Christopher K. Wyczalkowski, and Xi Huang. Public transit access and the changing spatial distribution of poverty. *Regional Science and Urban Economics*, 66:198–212, 2017.
- Steven Ruggles, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. Ipums usa: Version 9.0 [dataset]. Minneapolis, MN: IPUMS, 2019.
- Chris Severen. Commuting, labor, and housing market effects of mass transportation: Welfare and identification. Federal Reserve Bank of Philadelphia Working Paper, 2019.
- Kenneth A. Small. Valuation of travel time. *Economics of Transportation*, 1(1):2–14, 2012.
- Nick Tsivanidis. Evaluating the impact of urban transit infrastructure: Evidence from bogotá’s transmilenio. Working Paper, 2019.
- Federal Highway Administration U.S. Department of Transportation. 2017 national household travel survey. URL: <http://nhts.ornl.gov>.
- Pablo Ernesto Warnes. Transport infrastructure improvements and spatial sorting: Evidence from buenos aires. Working paper, 2020.
- Román David Zárate. Spatial misallocation, informality and transit improvements: Evidence from mexico city. Working Paper, 2020.

A Appendix

A.1 Defining and simulating trips on Google Maps

This section describes how I sample trip origin-destination pairs and query them on Google Maps for travel times. These travel times are then used to estimate census tract-specific travel speeds by each mode as described in Section A.2. To do so credibly, there needs to be enough trips going through each census tract and the speeds on these trips need to be representative of trips actually taken (and not just speeds on infrequently travelled routes). In addition, because transit networks can be sparse, the trips need to be geographically spread out so that I am not just looking at areas within a tract far away from (or close to) transit routes.

To that end, a subset of trips are defined to be between the origins and destinations on trips reported in the 2017 National Household Travel Survey (NHTS). The confidential version of the NHTS (U.S. Department of Transportation) identifies locations at the block group level, and I define my trips to be between the population-weighted centroids of these block groups. I ignore round trips and unrepresentative trips (such as trips by air). NHTS trips span only a few thousand in total across all cities and they are missing for a sizeable share of my urban block groups. So, I generate additional origin-destination pairs myself.

For the remainder of my trips, I set the trip origins to be the population-weighted centroids of each block group with a non-zero residential population within the extent of my CBSAs. Since block groups are geographically smaller than tracts, I always guarantee a few trips originating in every tract. Trip destinations are of two types: (1) centroids of tracts that are popular commuting destinations as observed in the CTPP and (2) popular non-residential amenities nearby (such as restaurants and shopping malls). Popular commuting tract destinations include the 5 most popular destinations from the trip’s tract of origin and the 5 most popular destinations from the trip’s county of origin.

For trips to amenities, I first gather a dataset of popular amenities (also from Google). I categorize non-residential amenities into 19 types, each corresponding to a different Google “place type” on Google’s Places API. These amenity types include banks, cafes, churches, city hall, convenience stores, doctors, gyms, hospitals, libraries, mosques, movie theaters, parks, pharmacies, post offices, restaurants, schools, shopping malls, stadiums, and train stations. I use each of these as search terms to query

Google’s Places API for the most popular destinations of each type within a fixed radius of (the centroid of) each block group. On any search Google returns upto 20 places in order of “prominence”, as determined “by a place’s ranking in Google’s index, global popularity, and other factors.” The search radius determines the average proximity to the returned destinations. I let the radius vary with the place type being searched since some types may be sparser across space than others (such as restaurants are typically more common than stadiums). In setting the search radii, I also try to mimic the distribution of trip distances observed in the NHTS. For each place, Google returns geographical coordinates (as well as other data not used in this paper), which I then use to define trips between each block group centroid and the closest of the (twenty) returned destinations around it. Sometimes, Google may find no destination of a particular type around a block group, in which case I choose the closest destination from the pool of places of that type returned on queries from other block groups in the city. In total, I defined roughly 2 million origin-destination (O-D) pairs across all cities to query on Google Maps for travel times. I defined more trips in cities with more block groups (as per the trip sampling strategy outlined above), which also translates to more trips in more populated cities.

All trips were queried on weekdays in the middle of June 2018. Google’s travel time predictions for driving and walking trips are based on its own historical and real-time data. I only scrape the travel times that are based on historical averages (as opposed to real-time predictions by Google). These averages do not vary much over time and should be less susceptible to idiosyncratic shocks at the time of the data collection. On the other hand, Google’s travel time predictions for transit are based on transit schedules shared by transit authorities and the GTFS. While the transit schedule variation is also small across weekdays, they are still sensitive to the trip’s departure time. So, I repeat each transit trip at roughly 5 different hours of the day and take a weighted average where the weights are constructed from the distribution of trip departure times observed in the NHTS.

Not all queries to Google Maps return route results. A small share of driving queries (less than 1%) and walking queries (less than 2%) return null results, but roughly a fifth of transit queries return null results. Transit networks are sparse, so this is unsurprising. In fact, the share of null results would be higher if not for Google returning the walking routes in most (but not all) of the cases where the trip does not overlap with any transit route. The rate of queries with non-zero returns varies across

and within cities, with more null results farther away from city centers. I impute travel times on missing trips by assuming people walk the entire trip (straight-line) distance at a speed that is the 90th percentile of 'effective' walking speeds across successful trips in the surrounding tracts. The 'effective' speed is the straight-line distance covered per minute and by penalizing the missing trips with a slow walking speed, I implicitly assume there are obstructions and long detours along the way (that are also leading Google to not return these as viable travel routes).

A.2 Estimating tract speeds and commuting times

The goal of this exercise is to estimate travel times by driving, transit and walking between all possible pairings of residential and work tracts within a city. I compute this matrix of travel times in three steps. First, I identify the shortest routes between all O-D pairs (including for the non-commuting trips queried on Google Maps) along major road networks and their overlap with the city's tracts. Second, using the trips for which I also have travel times from Google Maps, I estimate tract-level speeds on each travel mode using a series of OLS regressions. Third, I use the estimated speeds together with route overlaps with tracts to predict travel times on the remaining (commuting) trips.

I download the network of major streets in each city from OpenStreetMap (OSM), a crowd-sourced mapping platform. The street networks cover a 1% buffer zone around the geographic bounds of trips and includes the following OSM street types: motorways, trunks, primary, secondary, tertiary and unclassified. To improve the speed of (and memory constraints from) the millions of shortest route searches, I exclude smaller residential streets and driveways. As such, my 'shortest routes' are only along major streets and may not be the actual shortest route along the entire road network. This is not a major concern because residential streets tend to be slower and even when they make up a large portion of the shortest route, they are less likely to be part of the fastest route (or to be traveled). As such, my routes may even be more representative of actual traveled routes.

Then I map each trip origin and destination to their nearest point on the street network and project the entire network as a directional graph of edges along streets and nodes at street intersections and trip endpoints. Shortest paths between trip endpoints are computed using NetworkX, a python package. Having identified the shortest routes, I intersect them at tract boundaries and compute the lengths of the intersections with

each tract that is within a 1% buffer around the convex hull of the set of tracts in the CBSA. I ignore a small share of trips with less than 50% overlap with these tracts. Note that commuting trips are defined to be between the centroids of tracts within the CBSA and are, hence, always within the convex hull of these tracts. The total distance along the shortest path on commuting trips is the trip distance measure used in subsequent analysis from Section 3 onwards.

For estimating tract speeds, I specify a trip’s total travel time as the sum of travel times through each tract that its route overlaps. As shown in (2.1), I further decompose each travel time segment into a route distance divided by travel speed. When I know both the total travel time and the distances traveled through each tract, I can use an OLS regression to uncover the coefficients on distances which are also the travel speeds in the corresponding tracts. So, using my set of non-commuting trips for which I have the Google Maps travel times, I run separate regressions for each city and travel mode to estimate the tract-specific speeds.

With the large number of tracts to estimate speeds for, the OLS regression faces a multi-colinearity problem that is more prominent among tracts with limited variation in trip routes. For example, if two tracts share a large fraction of the trips passing through them, then it is difficult to isolate the effect that going through each tract has on the trips’ travel times. In the worst case, some tracts have to be dropped from the estimation due to perfect colinearity. I assign each dropped tract the median of estimated speeds of their surrounding tracts. The OLS regression may also estimate extremely high or low speeds for some less central tracts in the city. So, I truncate the top and bottom 5% of estimated speeds in each city.

Finally, to predict total travel times on commuting trips, I plug in the estimated tract speeds along with each trip’s route (length) overlaps with tracts into (2.1).

A.3 Housing demand estimation

I observe annual household incomes and annual housing expenditures in the publicly available census microdata from IPUMS but not their census tract of residence. The smallest identifiable geography of residence in the microdata are PUMAs, which are usually larger than tracts. In order to combine the housing expenditure data with tract-level standardized housing prices from Davis et al. [2020] to the microdata, I rely on aggregate tract-level data from NHGIS. In the aggregate data, I observe household counts by census tract across 16 income brackets, so I can determine the median housing

prices experienced by households in each bracket. To aggregate the tract level prices for each income group to the PUMA level, I use a crosswalk from the website of Missouri Census Data Center that returns the population-weighted overlaps between PUMAs and tracts. I use these as weights to compute the median housing prices experienced by households in each income bracket within a PUMA. I merge these PUMA-income-level average prices to the micro-data to use in the housing demand estimation. As such, these are not the actual housing prices corresponding to the housing expenditures but the housing prices likely to be experienced by the median household in the same income bracket and PUMA.

Once I have incomes, housing expenditures and housing prices for my sample of individual households, I run the OLS regression in (4.5) separately for each CBSA: I regress the log share of income spent on housing expenditures (on the left) on log income and log housing price (on the right). Observations are weighted by survey weights for households and excludes households with zero housing expenditure and households at the top and bottom percentiles of the sample’s income distribution. Sample sizes for the regressions range from around 15,000 households in the smallest cities to over 700,000 in the largest ones.

Estimated price elasticities of housing demand (α_h) range from -0.66 (in Syracuse, NY) to -0.82 (in San Francisco, CA). Estimated income elasticities of housing demand (α_w) range from 0.4 (in Provo-Orem, UT) to 0.6 (in San Francisco, CA). Figure A.1 compares the predicted (“fitted”) and the observed housing expenditures as a function of household income, pooling together all cities. My predicted housing expenditures are too high for household incomes below \$15,000 (an artifact of the log-log functional form) and slightly smaller than those observed for incomes above.

A.4 Mode and neighborhood choice estimation

Estimation requires numerically searching over parameters α_{my}^S and α_y^D and fixed effects δ_{mny} to maximize the sum of log likelihoods L from (4.6). To aid the search process, I exploit a contraction mapping approach popularized by Berry et al. [1995]. More specifically, given any realization of the vector of parameters α^S and α^D , a contraction mapping is used to calculate the matrix of fixed effects δ that solves the first order conditions $\frac{\partial L}{\partial \delta} = 0$.

Consider the following first-order Taylor approximation of L as a function of the




Figure A.1: **Predicted vs observed housing expenditures in the microdata.** The fitted regression line is based on the OLS estimation of (4.5). The figure pools together households across all cities.

fixed effects:

$$L(\alpha^S, \alpha^D, \delta^{t+1}) = L(\alpha^S, \alpha^D, \delta^t) + (\delta^{t+1} - \delta^t) \frac{\partial L(\alpha^S, \alpha^D, \delta^t)}{\partial \delta}$$

The first order condition to solve for the δ^{t+1} that maximizes this approximation is

$$\frac{\partial L(\delta^{t+1})}{\partial \delta^{t+1}} = 0$$

which, following some algebraic manipulation, evaluates to

$$\delta_{mny}^{t+1} = \delta_{mny}^t - \ln \left[\frac{\sum_j \left(\sum_m \sum_n \sum_y P_{jmny} \right) \pi_{mn|jy}}{\left(\sum_j P_{jmny} \right)} \right]$$

Updating the values of δ as above until convergence maximizes L conditional on parameters α^S and α^D . I update parameters α^S and α^D by the (weighted) gradient of

the log likelihood with respect to each:

$$\frac{\partial L}{\partial \alpha_y^D} = \sum_j \sum_n \sum_{m \in M} \left[P_{jmny} - \pi_{mn|jy} \left(\sum_m \sum_n P_{jmny} \right) \right] \cdot D_{jn}$$

$$\frac{\partial L}{\partial \alpha_{my}^S} = \sum_j \sum_n \left[P_{jmny} - \pi_{mn|jy} \left(\sum_m \sum_n P_{jmny} \right) \right] \cdot S_{jmn}$$

A.5 Additional tables and figures

Travel distances vs speeds

Figure A.2 shows that average travel distance is increasing with travel speed, suggesting commuters (on average) travel farther when they can travel faster. Alternatively, longer trips tend to be faster. This is true for both drivers and transit riders and across all income groups. However, even conditional on speed, higher income groups appear to commute slightly longer..

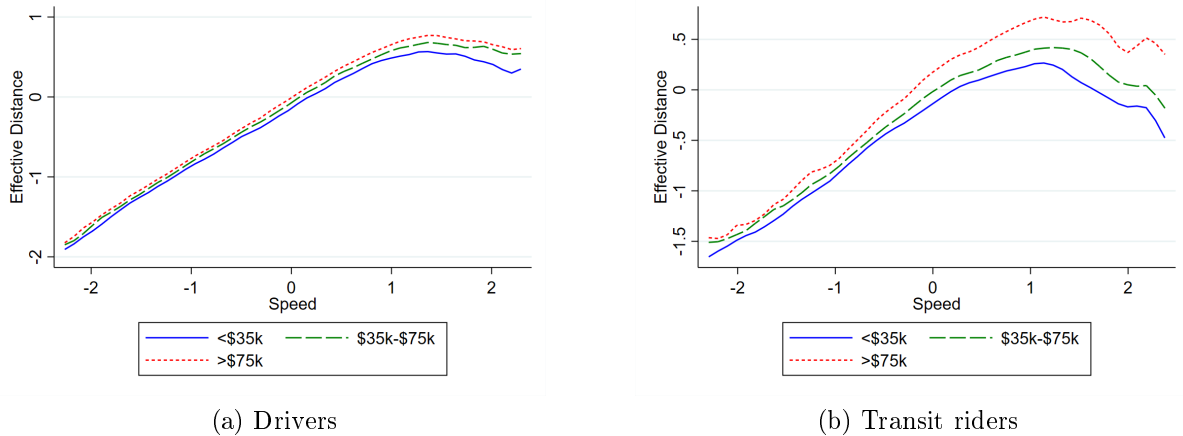


Figure A.2: **Average commuting distances by speed.** Distances and speeds are in logs and standardized across commutes within each CBSA and travel mode. The figures pool together commutes across all cities in my sample.

Table A.1: Full ranking of cities by commuting speeds on transit

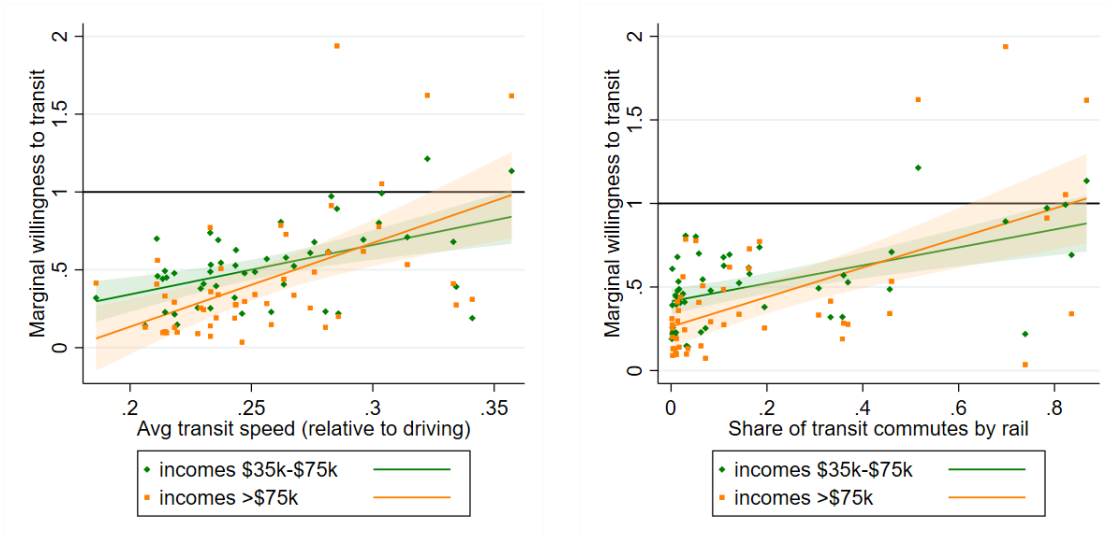
Rank	City	Transit speed (in km/h)	Ratio of transit to driving speed	% commuters riding transit	Rail share of transit riders
1	New York-Newark-Jersey City, NY-NJ-PA	20.2	0.35	30.7%	86.7%
2	San Francisco-Oakland-Hayward, CA	18.9	0.31	15.5%	51.6%
3	Seattle-Tacoma-Bellevue, WA	17.8	0.30	8.6%	5.2%
4	Chicago-Naperville-Elgin, IL-IN-WI	17.7	0.28	11.9%	69.8%
5	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	17.1	0.30	9.7%	46.0%
6	Washington-Arlington-Alexandria, DC-VA-MD-WV	17.0	0.30	14.5%	82.3%
7	Houston-The Woodlands-Sugar Land, TX	16.5	0.21	2.7%	2.5%
8	Denver-Aurora-Lakewood, CO	15.9	0.28	4.9%	16.2%
9	Urban Honolulu, HI	15.8	0.33	8.2%	0.3%
10	Atlanta-Sandy Springs-Roswell, GA	15.7	0.21	3.5%	30.8%
11	Santa Maria-Santa Barbara, CA	15.7	0.28	4.0%	1.3%
12	Los Angeles-Long Beach-Anaheim, CA	15.5	0.27	6.4%	10.9%
13	Boston-Cambridge-Newton, MA-NH	15.5	0.28	12.4%	78.4%
14	St. Louis, MO-IL	15.4	0.23	2.7%	19.5%
15	Minneapolis-St. Paul-Bloomington, MN-WI	15.1	0.20	4.8%	5.8%
16	Portland-Vancouver-Hillsboro, OR-WA	14.9	0.29	6.6%	12.2%
17	Sacramento-Roseville-Arden-Arcade, CA	14.8	0.23	2.8%	18.4%
18	Miami-Fort Lauderdale-West Palm Beach, FL	14.7	0.26	3.8%	14.2%
19	Cleveland-Elyria, OH	14.6	0.21	4.0%	8.3%
20	San Antonio-New Braunfels, TX	14.5	0.23	2.3%	0.3%
21	Milwaukee-Waukesha-West Allis, WI	14.4	0.23	3.8%	1.5%
22	Phoenix-Mesa-Scottsdale, AZ	14.4	0.23	2.3%	2.8%
23	Pittsburgh, PA	14.3	0.25	6.0%	3.0%
24	Boulder, CO	14.2	0.32	5.9%	0.2%
25	Hartford-West Hartford-East Hartford, CT	14.0	0.23	2.9%	6.6%
26	San Diego-Carlsbad, CA	13.9	0.24	3.5%	11.0%
27	Salt Lake City, UT	13.9	0.26	3.5%	16.3%
28	Providence-Warwick, RI-MA	13.8	0.24	2.7%	45.6%
29	Baltimore-Columbia-Towson, MD	13.7	0.25	6.5%	36.0%
30	Eugene, OR	13.7	0.28	4.1%	0.3%
31	Albany-Schenectady-Troy, NY	13.4	0.25	3.2%	6.2%
32	Las Vegas-Henderson-Paradise, NV	13.3	0.27	3.8%	0.2%
33	Buffalo-Cheektowaga-Niagara Falls, NY	13.1	0.23	3.8%	7.2%
34	Bridgeport-Stamford-Norwalk, CT	13.1	0.23	9.7%	83.6%
35	Rochester, NY	13.1	0.19	2.1%	3.5%
36	Madison, WI	13.0	0.23	4.0%	1.1%
37	Austin-Round Rock, TX	12.8	0.21	2.8%	1.1%
38	Tucson, AZ	12.8	0.28	2.6%	0.4%
39	Durham-Chapel Hill, NC	12.6	0.24	4.0%	1.4%
40	Lansing-East Lansing, MI	12.5	0.21	2.4%	0.8%
41	Provo-Orem, UT	12.5	0.26	2.2%	2.0%
42	Syracuse, NY	11.8	0.21	2.2%	1.0%
43	San Jose-Sunnyvale-Santa Clara, CA	11.7	0.24	3.4%	36.9%
44	Trenton, NJ	11.1	0.24	7.6%	73.9%
45	New Haven-Milford, CT	11.1	0.24	3.8%	35.8%
46	Ann Arbor, MI	10.8	0.23	4.1%	1.7%
47	Springfield, MA	10.7	0.21	2.1%	3.2%
48	Savannah, GA	9.9	0.21	2.1%	0.9%
49	Vallejo-Fairfield, CA	8.8	0.18	2.7%	33.3%

Note: Speeds are relative to shortest road distance. Speeds and ratios of travel times are means across all trips between observed work-residence location pairs (unconditional on travel mode choice) ignoring the top and bottom 5% of outliers. Rail share is the fraction of transit commutes via rail transit in the city.

Table A.2: Estimated coefficients α_y^D and α_{my}^S on commuting distance and speed

Variable	Mode	Income	Mean	p5	p25	Median	p75	p95
Distance	all	< \$35k	1.518	0.971	1.260	1.601	1.686	2.006
		\$35k-\$50k	1.501	0.911	1.260	1.555	1.700	1.991
		\$50k-\$75k	1.453	0.864	1.142	1.514	1.649	1.947
		> \$75k	1.421	0.851	1.155	1.465	1.624	1.891
Speed	driving	< \$35k	0.301	0.156	0.218	0.293	0.370	0.471
		\$35k-\$50k	0.310	0.131	0.247	0.303	0.376	0.494
		\$50k-\$75k	0.293	0.083	0.211	0.297	0.380	0.487
		> \$75k	0.272	0.091	0.201	0.273	0.337	0.484
Speed	transit	< \$35k	0.271	0.140	0.210	0.271	0.313	0.387
		\$35k-\$50k	0.241	0.119	0.197	0.220	0.257	0.374
		\$50k-\$75k	0.243	0.087	0.200	0.248	0.282	0.377
		> \$75k	0.276	0.097	0.225	0.274	0.319	0.478
Speed	walking	< \$35k	-0.155	-0.303	-0.212	-0.153	-0.103	0.010
		\$35k-\$50k	-0.146	-0.364	-0.189	-0.146	-0.086	0.054
		\$50k-\$75k	-0.177	-0.408	-0.219	-0.148	-0.115	-0.039
		> \$75k	-0.265	-0.435	-0.321	-0.265	-0.208	-0.137

Note: Table reports the mean, 5th percentile, 25th percentile, median, 75th percentile and 95th percentile (in that order) of coefficient estimates across all 49 cities. Standard errors on all estimated coefficients are less than 0.00001.



(a) by mean transit speeds (relative to driving) (b) by rail share of transit commutes in the city

Figure A.3: Mean MWT from 1% increase in transit speed (relative to lowest income group). Each observation corresponds to a city. Vertical axis depicts the MWT for faster transit as a fraction of the MWT of commuters with incomes less than \$35,000 (indicated by solid black line at 1). Horizontal axis depicts either (a) the ratio of driving to transit travel times (across all observed commutes) in the city or (b) the share of transit riders in the city who commute by rail transit. Confidence intervals for each linear fit are shaded in corresponding color. For commuters with incomes \$35k-\$75k, figures plot population-weighted means of the MWT estimates for the two middle-income groups in my data.

Table A.3: Cities ranked by mean MWTP for faster transit commutes

Rank	City	MWTP for faster transit	MWTP for faster driving
1	San Francisco-Oakland-Hayward, CA	\$ 374	\$ 302
2	Seattle-Tacoma-Bellevue, WA	\$ 188	\$ 179
3	New York-Newark-Jersey City, NY-NJ-PA	\$ 178	\$ 345
4	San Jose-Sunnyvale-Santa Clara, CA	\$ 169	\$ 139
5	Boston-Cambridge-Newton, MA-NH	\$ 148	\$ 189
6	Washington-Arlington-Alexandria, DC-VA-MD-WV	\$ 129	\$ 156
7	Vallejo-Fairfield, CA	\$ 119	\$ 69
8	Chicago-Naperville-Elgin, IL-IN-WI	\$ 116	\$ 179
9	Los Angeles-Long Beach-Anaheim, CA	\$ 114	\$ 102
10	Sacramento-Roseville-Arden-Arcade, CA	\$ 101	\$ 95
11	Portland-Vancouver-Hillsboro, OR-WA	\$ 97	\$ 94
12	Denver-Aurora-Lakewood, CO	\$ 94	\$ 67
13	San Diego-Carlsbad, CA	\$ 87	\$ 92
14	Boulder, CO	\$ 85	\$ 48
15	Minneapolis-St. Paul-Bloomington, MN-WI	\$ 78	\$ 129
16	Providence-Warwick, RI-MA	\$ 78	\$ 134
17	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	\$ 77	\$ 123
18	Pittsburgh, PA	\$ 73	\$ 90
19	Houston-The Woodlands-Sugar Land, TX	\$ 70	\$ 70
20	Miami-Fort Lauderdale-West Palm Beach, FL	\$ 64	\$ 75
21	Bridgeport-Stamford-Norwalk, CT	\$ 64	\$ 80
22	Hartford-West Hartford-East Hartford, CT	\$ 61	\$ 74
23	Santa Maria-Santa Barbara, CA	\$ 61	\$ 78
24	Atlanta-Sandy Springs-Roswell, GA	\$ 59	\$ 78
25	Baltimore-Columbia-Towson, MD	\$ 59	\$ 76
26	Savannah, GA	\$ 52	\$ 19
27	New Haven-Milford, CT	\$ 47	\$ 54
28	San Antonio-New Braunfels, TX	\$ 47	\$ 34
29	Phoenix-Mesa-Scottsdale, AZ	\$ 44	\$ 47
30	St. Louis, MO-IL	\$ 43	\$ 47
31	Ann Arbor, MI	\$ 42	\$ 27
32	Albany-Schenectady-Troy, NY	\$ 41	\$ 55
33	Austin-Round Rock, TX	\$ 39	\$ 65
34	Eugene, OR	\$ 38	\$ 26
35	Milwaukee-Waukesha-West Allis, WI	\$ 35	\$ 39
36	Urban Honolulu, HI	\$ 32	\$ 19
37	Durham-Chapel Hill, NC	\$ 32	\$ 56
38	Springfield, MA	\$ 32	\$ 43
39	Madison, WI	\$ 31	\$ 42
40	Trenton, NJ	\$ 31	\$ 45
41	Lansing-East Lansing, MI	\$ 30	\$ 30
42	Cleveland-Elyria, OH	\$ 30	\$ 57
43	Tucson, AZ	\$ 29	\$ 35
44	Salt Lake City, UT	\$ 28	\$ 19
45	Syracuse, NY	\$ 26	\$ 38
46	Provo-Orem, UT	\$ 22	\$ 17
47	Rochester, NY	\$ 21	\$ 50
48	Buffalo-Cheektowaga-Niagara Falls, NY	\$ 17	\$ 36
49	Las Vegas-Henderson-Paradise, NV	\$ 9	\$ 17

Note: Cities are ranked by the mean MWTP for faster transit. MWTP values are means across all commuters for 1% change in travel speed on their observed commutes (i.e. conditional on commuters choosing their observed modes and neighborhoods).

Table A.4: Cities ranked by MWTT from 1% increase in commuting speed


Rank	City	%age pt change in transit ridership	Baseline transit ridership (in %)
1	San Francisco-Oakland-Hayward, CA	3.52	15.45
2	Chicago-Naperville-Elgin, IL-IN-WI	2.97	11.87
3	Washington-Arlington-Alexandria, DC-VA-MD-WV	2.97	14.46
4	Seattle-Tacoma-Bellevue, WA	2.88	8.63
5	Boston-Cambridge-Newton, MA-NH	2.76	12.43
6	Portland-Vancouver-Hillsboro, OR-WA	2.22	6.57
7	New York-Newark-Jersey City, NY-NJ-PA	2.22	30.74
8	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	1.81	9.66
9	Pittsburgh, PA	1.77	5.97
10	Boulder, CO	1.73	5.93
11	Denver-Aurora-Lakewood, CO	1.57	4.92
12	Baltimore-Columbia-Towson, MD	1.16	6.51
13	Los Angeles-Long Beach-Anaheim, CA	1.11	6.38
14	Minneapolis-St. Paul-Bloomington, MN-WI	1.07	4.76
15	Eugene, OR	1.05	4.11
16	Atlanta-Sandy Springs-Roswell, GA	0.97	3.54
17	Houston-The Woodlands-Sugar Land, TX	0.85	2.66
18	Miami-Fort Lauderdale-West Palm Beach, FL	0.80	3.84
19	Durham-Chapel Hill, NC	0.79	4.02
20	Savannah, GA	0.79	2.14
21	Madison, WI	0.78	4.01
22	Bridgeport-Stamford-Norwalk, CT	0.76	9.73
23	San Antonio-New Braunfels, TX	0.76	2.32
24	Salt Lake City, UT	0.75	3.46
25	Urban Honolulu, HI	0.71	8.18
26	Ann Arbor, MI	0.70	4.06
27	New Haven-Milford, CT	0.69	3.78
28	Sacramento-Roseville-Arden-Arcade, CA	0.69	2.83
29	San Jose-Sunnyvale-Santa Clara, CA	0.66	3.44
30	Tucson, AZ	0.65	2.63
31	Santa Maria-Santa Barbara, CA	0.56	4.01
32	Austin-Round Rock, TX	0.55	2.78
33	St. Louis, MO-IL	0.55	2.69
34	Phoenix-Mesa-Scottsdale, AZ	0.55	2.33
35	San Diego-Carlsbad, CA	0.54	3.52
36	Albany-Schenectady-Troy, NY	0.51	3.24
37	Hartford-West Hartford-East Hartford, CT	0.45	2.86
38	Lansing-East Lansing, MI	0.45	2.40
39	Trenton, NJ	0.39	7.57
40	Vallejo-Fairfield, CA	0.38	2.71
41	Buffalo-Cheektowaga-Niagara Falls, NY	0.38	3.77
42	Provo-Orem, UT	0.36	2.22
43	Milwaukee-Waukesha-West Allis, WI	0.35	3.77
44	Cleveland-Elyria, OH	0.34	4.03
45	Providence-Warwick, RI-MA	0.33	2.72
46	Springfield, MA	0.33	2.08
47	Syracuse, NY	0.26	2.16
48	Las Vegas-Henderson-Paradise, NV	0.21	3.75
49	Rochester, NY	0.13	2.08

Note: Cities are ranked by the MWTT in response to a 1% increase in transit speed along all observed commutes (i.e. conditional on commuters choosing their observed neighborhoods).


Table A.5: Cities ranked by unconditional MWTP for faster transit

Rank	City	uMWTP for faster transit			Relative to driving	
		all commuters	incomes <\$35k	incomes >\$75k	all commuters	com-muters
1	San Francisco-Oakland-Hayward, CA	\$ 39.82	\$ 17.24	\$ 46.26	0.1179	
2	New York-Newark-Jersey City, NY-NJ-PA	\$ 38.75	\$ 18.22	\$ 44.21	0.0954	
3	Boston-Cambridge-Newton, MA-NH	\$ 16.19	\$ 12.90	\$ 16.04	0.0750	
4	Washington-Arlington-Alexandria, DC-VA-MD-WV	\$ 14.81	\$ 13.86	\$ 14.31	0.0868	
5	Chicago-Naperville-Elgin, IL-IN-WI	\$ 11.44	\$ 4.81	\$ 16.29	0.0534	
6	Seattle-Tacoma-Bellevue, WA	\$ 7.87	\$ 9.33	\$ 6.96	0.0379	
7	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	\$ 5.69	\$ 8.20	\$ 4.49	0.0379	
8	Portland-Vancouver-Hillsboro, OR-WA	\$ 2.49	\$ 4.26	\$ 1.67	0.0230	
9	Pittsburgh, PA	\$ 2.01	\$ 3.51	\$ 1.52	0.0181	
10	Minneapolis-St. Paul-Bloomington, MN-WI	\$ 1.59	\$ 3.38	\$ 0.92	0.0103	
11	Los Angeles-Long Beach-Anaheim, CA	\$ 1.56	\$ 2.91	\$ 0.89	0.0121	
12	Baltimore-Columbia-Towson, MD	\$ 1.37	\$ 4.11	\$ 0.55	0.0158	
13	Atlanta-Sandy Springs-Roswell, GA	\$ 0.97	\$ 2.37	\$ 0.51	0.0101	
14	Denver-Aurora-Lakewood, CO	\$ 0.79	\$ 1.75	\$ 0.37	0.0104	
15	Miami-Fort Lauderdale-West Palm Beach, FL	\$ 0.72	\$ 1.51	\$ 0.37	0.0078	
16	Boulder, CO	\$ 0.60	\$ 2.75	\$ 0.05	0.0120	
17	San Jose-Sunnyvale-Santa Clara, CA	\$ 0.60	\$ 2.10	\$ 0.15	0.0039	
18	Santa Maria-Santa Barbara, CA	\$ 0.56	\$ 1.19	\$ 0.26	0.0064	
19	Urban Honolulu, HI	\$ 0.54	\$ 3.13	\$ 0.24	0.0253	
20	Durham-Chapel Hill, NC	\$ 0.52	\$ 1.41	\$ 0.11	0.0084	
21	San Diego-Carlsbad, CA	\$ 0.51	\$ 1.37	\$ 0.18	0.0049	
22	Ann Arbor, MI	\$ 0.50	\$ 1.42	\$ 0.02	0.0179	
23	Bridgeport-Stamford-Norwalk, CT	\$ 0.50	\$ 1.52	\$ 0.18	0.0054	
24	Sacramento-Roseville-Arden-Arcade, CA	\$ 0.46	\$ 0.90	\$ 0.27	0.0042	
25	Albany-Schenectady-Troy, NY	\$ 0.39	\$ 2.38	\$ 0.04	0.0063	
26	Austin-Round Rock, TX	\$ 0.38	\$ 1.24	\$ 0.05	0.0052	
27	Providence-Warwick, RI-MA	\$ 0.35	\$ 0.87	\$ 0.18	0.0022	
28	Madison, WI	\$ 0.35	\$ 1.32	\$ 0.05	0.0073	
29	Houston-The Woodlands-Sugar Land, TX	\$ 0.35	\$ 0.74	\$ 0.22	0.0041	
30	New Haven-Milford, CT	\$ 0.34	\$ 1.54	\$ 0.06	0.0056	
31	Hartford-West Hartford-East Hartford, CT	\$ 0.28	\$ 0.86	\$ 0.17	0.0032	
32	Springfield, MA	\$ 0.27	\$ 1.41	\$ 0.02	0.0056	
33	Trenton, NJ	\$ 0.26	\$ 1.93	\$ 0.00	0.0054	
34	Eugene, OR	\$ 0.25	\$ 1.05	\$ 0.02	0.0090	
35	Phoenix-Mesa-Scottsdale, AZ	\$ 0.24	\$ 0.65	\$ 0.10	0.0044	
36	Cleveland-Elyria, OH	\$ 0.22	\$ 0.79	\$ 0.07	0.0032	
37	St. Louis, MO-IL	\$ 0.21	\$ 0.76	\$ 0.07	0.0038	
38	Buffalo-Cheektowaga-Niagara Falls, NY	\$ 0.21	\$ 0.97	\$ 0.02	0.0051	
39	San Antonio-New Braunfels, TX	\$ 0.17	\$ 0.51	\$ 0.01	0.0044	
40	Lansing-East Lansing, MI	\$ 0.14	\$ 0.26	\$ 0.14	0.0044	
41	Salt Lake City, UT	\$ 0.09	\$ 0.35	\$ 0.03	0.0046	
42	Savannah, GA	\$ 0.09	\$ 0.23	\$ 0.00	0.0045	
43	Rochester, NY	\$ 0.09	\$ 0.59	\$ 0.01	0.0015	
44	Tucson, AZ	\$ 0.08	\$ 0.24	\$ 0.01	0.0022	
45	Vallejo-Fairfield, CA	\$ 0.06	\$ 0.14	\$ 0.01	0.0008	
46	Syracuse, NY	\$ 0.04	\$ 0.20	\$ 0.00	0.0009	
47	Milwaukee-Waukesha-West Allis, WI	\$ 0.03	\$ 0.13	\$ 0.01	0.0008	
48	Las Vegas-Henderson-Paradise, NV	\$ 0.03	\$ 0.12	\$ 0.00	0.0019	
49	Provo-Orem, UT	\$ 0.02	\$ 0.06	\$ 0.00	0.0012	


Note: Cities are ranked by their mean uMWTP (across all commuters) for 1% increase in transit speeds. Reported uMWTP values are estimates of mean MWTP across all commuters unconditional on their choices of mode and neighborhood. Ratio in column 6 divides uMWTP estimates in column 3 by estimates of the city's mean uMWTP for 1% increase in driving speeds.




(a) Seattle, Incomes < \$35k



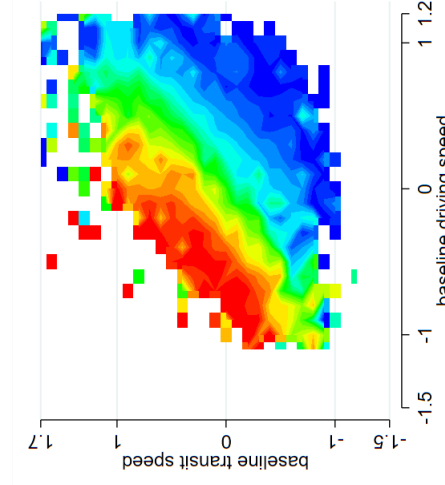
(b) Seattle, Incomes \$35k-\$75k




(c) Seattle, Incomes > \$75k



(d) Boston, Incomes < \$35k



(e) Boston, Incomes \$35k-\$75k



(f) Boston, Incomes > \$75k

Figure A.4: Mean MWT by location of transit improvement. The x- and y-axes depict standardized driving and transit speeds (resp.) on the commuting route. The z-axis depicts the percentage point change in transit ridership in response to a 1% increase in transit speed along the commuting route. The z-axis colours are fixed across all graphs. Speeds are standardized (to mean 0 and std. dev. 1) across trips between observed work-residence pairs within each CBSA.. Trips at the top and bottom percentiles of speeds are ignored. White spaces in the graphs correspond to 0.1-by-0.1 cells with fewer than 20 commutes.