

Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio*

Nick Tsivanidis[†]

University of California, Berkeley

August, 2023

Abstract

This paper estimates the effects of improving public transit infrastructure on city structure and welfare. It begins by developing a quantitative urban model with multiple groups of workers and transit modes. A special case of this model admits a sufficient statistics approach that measures aggregate welfare gains from improved transit in a broader class of models. The paper then estimates the reduced-form elasticities needed to implement the approach using data spanning the construction of the world's largest Bus Rapid Transit system in Bogotá, Colombia. This class of models performs well in explaining the adjustment of economic activity to the system. The standard approach for measuring the welfare gains from new infrastructure based on the value of travel time saved only accounts for 54% of the total welfare gain. Using the more general model to assess the distributional consequences, there is little impact on inequality after accounting for reallocation and general equilibrium effects.

*Thanks to Chang-Tai Hsieh, Marianne Bertrand, Erik Hurst and Robert E. Lucas Jr. for their guidance and support. I also thank Rodrigo Adao, Treb Allen, David Atkin, Jonathan Dingel, Cecile Gaubert, Rick Hornbeck, Daniel Sturm, Felix Tintelnot, Owen Zidar, and numerous seminar participants for their helpful comments. Additionally, I am thankful to Andres Villaveces, Marco Antonio Llinas Vargas, Helmut Menjura and those at DANE for help with data access. I gratefully acknowledge financial support from the Chicago Booth Social Enterprise Initiative, the University of Chicago Urban Network, the Energy Policy Institute at Chicago and the North American Regional Science Council. A previous version of this paper circulated under the title "The Aggregate and Distributional Effect of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio". The results in this paper use confidential microdata from DANE and have been reviewed to ensure no confidential information has been disclosed. The opinions and conclusions expressed herein are my own and do not necessarily represent the views of DANE. All errors are mine.

[†]Haas School of Business, University of California at Berkeley. Email: ntsivanidis@berkeley.edu. First version: November 2017.

1 Introduction

How large are the economic gains from improving public transit systems in cities? With 2.5 billion people predicted to move mostly into developing country cities by 2050, governments will spend vast sums on mass transit to reduce the congestion associated with rapid urban growth.¹ While existing approaches focus on the value of travel time saved (VTTS), measuring the benefits of these systems is challenging.² Individuals may change where they live and work, firms may expand or enter in newly accessible locations, and wages and house prices may adjust to this reallocation. Such effects are missed by time savings, and indirect effects may be felt throughout the city even on those who do not use the system. The lack of detailed intra-city data in less developed countries coinciding with the opening of large transit systems makes the task of evaluating these impacts even more daunting.

This study investigates the impact of new transit infrastructure on the structure of cities and the welfare of their inhabitants. It does so in the context of the construction of the world's most-used Bus Rapid Transit (BRT) system—TransMilenio—in Bogotá, Colombia. Opened in 2000, TransMilenio has a daily volume of over 2.2 million trips and operates similar to a subway. Buses run in dedicated lanes with express and local services, and passengers board buses at stations which they pay to enter using smart cards. BRT provides an attractive alternative to subways in rapidly growing developing country cities: they can deliver similar reductions in commuting times at a fraction of the cost, and are much faster to build.³ This paper uses new sources of data covering 2,800 census tracts on residence, employment, commuting patterns, and land markets spanning the system's construction.

Prior to TransMilenio, poor and low-educated workers relied on a network of informal buses which, on average, were 30% slower than cars. This suggests new transit may affect the distribution of welfare across the rich and poor. To better understand the implications of improving transit infrastructure, this paper develops a quantitative urban model. Multiple worker skill groups have non-homothetic preferences over transit modes and residential locations, and make decisions over where to live and work and which travel mode to use to commute. Cars are faster than public transit, but are expensive. In equilibrium, rich, high-educated workers are more likely to buy cars while the poor rely on public transit. While this suggests the poor are likely to benefit most from improved public transit, worker types differ in their willingness to substitute between alternative residential and employment locations and are exposed to equilibrium effects on wages and house prices.

While this model is rich enough to speak to distributional impacts, a special case admits an even simpler sufficient statistics approach to measure the aggregate from new transit infrastructure and how it it reshapes economic activity across the city. This approach has appeal since these statistics are transparently estimated through linear regression, and because the approach is applicable in a

¹McKinsey (2016) suggests that a need for \$40 trillion of spending to close the transport infrastructure gap. Combining the average subway distance from Gonzalez-Navarro and Turner (2018) and cost estimates from Baum-Snow and Kahn (2005) indicates that the average subway system costs \$27.81bn in 2017 dollars to build.

²E.g. Train and McFadden (1978), Small and Verhoef (2007), also used the World Bank (Mackie et. al. 2005)

³The per mile construction cost of the subway in Colombia's second largest city, Medellín, was 10 times that of TransMilenio, with similar system speeds. TransMilenio took less than 18 months to construct, compared to the 12 years taken by Metro Medellín. The average per mile construction cost of BRT is one-tenth of rail (Menckhoff 2005).

wide class of log-linear models that allow for endogenous firm location choice, endogenous housing supply, capital in the production function, and preferences over leisure (among others). These statistics are (i) a location's change in "commuter market access" (CMA), which summarizes worker and firm access to each other through the commuting network, and (ii) the elasticities of residential population, employment and floorspace prices to CMA, and the elasticity of commute flows to commute costs. Across alternative models, the structural parameters in the reduced-form elasticities of economic activity to CMA differ, but the reduced-form elasticities and the change in CMA are sufficient statistics to specify the impacts of changes in transit infrastructure on economic activity.⁴

The construction of TransMilenio provides variation in commute costs that can be used to estimate these elasticities, but concerns remain that these were endogenous to local unobserved economic fundamentals. Instead of leaning on a single approach, this paper exploits a variety of TransMilenio's institutional features to establish its causal impact on Bogotá's structure. First, I digitize four different plans from the 1980s and 1990s for a new transit network in Bogotá and include as regressors both the realized change in CMA due to TransMilenio and the hypothetical change had the network been built under these plans. This serves both as a falsification check (by showing the hypothetical changes had no impact on economic activity conditional on the realized CMA change) and controls for the omitted variable bias that can arise from locations' non-random exposure to infrastructure changes (Borusyak and Hull 2021). Second, I exploit TransMilenio's staggered rollout across three phases through event studies and falsification tests and demonstrate that there is no growth in outcomes prior to line openings. Third, I use variation in CMA induced by changes in the network more than 1.5km from a location, which is less likely to be correlated with local unobservables. Fourth, I condition on distance to the closest station to assess whether effects are driven by changes in accessibility rather than by other features of stations (e.g. changes in foot traffic or pollution). Fifth, I construct cost-shifting instruments to predict TransMilenio's routes based both on a historical tram network and engineering estimates of the cost to build BRT on different types of land.

After showing that the log-linear relationships between changes in outcomes and CMA predicted by this class of models are borne out in the data, I use the sufficient statistics approach to quantify the aggregate effects of the new infrastructure. A key theoretical result that arises through the application of the envelope theorem to the social planner's problem in an efficient economy is that the elasticity of welfare to a change in transit infrastructure is proportional to a weighted average of time savings. This is precisely the VTTS expression used in the literature: when the equilibrium is efficient and the change in infrastructure is infinitesimally small, only the direct effects of time saved matter. However, my results show that the VTTS only accounts for 54% of the total welfare gains under the equilibrium model. The size of the shock explains one-third of the gap and the externalities explain two-thirds. Welfare rose by 2.28% in the baseline case where the BRT does not cause migration into Bogotá from the rest of Colombia, and 0.6% with migration. GDP per capita rose by 2.5-5% in these cases respectively, net of construction and operating costs. Overall, TransMilenio can account for

⁴More precisely, as Proposition 1 establishes, the reduced form elasticities and changes in CMA are sufficient statistics across all such models to compute the relative change in economic activity in any location. To pin down the overall level of changes, one or two additional parameters are usually needed. These differ by model and are often readily calibrated.

between 2.96-14.36% of GDP growth in Bogotá from 2000 to 2016, and up to 29.24% of observed population growth. While these findings are specific to Bogotá, the framework can be applied to other cities in both developed and developing countries. While I do not find evidence TransMilenio impacted travel times on other modes, an extension of the model allowing for traffic congestion leaves these results qualitatively unchanged.

I next estimate the full model to understand how these welfare gains are shared between the rich and poor. Ultimately, welfare inequality rises by a mild 0.55% due to the BRT. On the one hand, low-skilled workers benefit from improved transit through higher use. On the other hand, the incidence also depends on how easily individuals substitute between different employment and residential locations and the extent to which each group faces lower wages through the increased supply of commuters traveling through the network. These forces favor high-skilled workers. This result is robust to allowing for employment in domestic services and alternative home ownership assumptions.

Two sets of counterfactuals draw additional policy insights. First, I evaluate a “land value capture” (LVC) scheme under which development rights to increase building densities near stations are sold by the government to developers. This increases housing supply and raises government revenue, and similar schemes have seen great success in Asian cities like Hong Kong and Tokyo. However, one of the main criticisms of TransMilenio was that the city experienced such a large change in transit without any adjustment to zoning laws to allow housing supply to respond. A well-targeted scheme would have increased the welfare gains from TransMilenio by around 44%, while government revenues would have covered 6-23% of the BRT’s capital costs depending on the migration response from the rest of Colombia. This highlights the return to cities pursuing an integrated transit and land use policy. Second, by measuring the impacts of counterfactual networks I find that the system of feeder buses, which run on regular roads and connect dense, outlying residential neighborhoods with TransMilenio terminals, have greater welfare gains than either of the two key trunk lines (conditional on the rest of the network being built). This emphasizes the importance of cheap, last-mile services that increase access to mass rapid transit infrastructure.

A large body of work examines the impact of transportation infrastructure on economic activity. One strand examines the impact of new transit infrastructure and typically measures changes in population and property prices as a function of distance to the central business district (Baum-Snow 2007; Gonzalez-Navarro and Turner 2018; Baum-Snow et. al. 2017) or distance to stations (Gibbons and Machin 2005; Glaeser et. al. 2008; Billings 2011). This paper adds to this work by developing a theory-consistent sufficient statistics approach to measure the impacts of transit infrastructure. The CMA measures used in this approach embrace the spillovers across spatial units induced through a commuting network that can invalidate identification assumptions in distance-based regressions.⁵

A second strand of this literature explores the effect of infrastructure between regions on economic development through goods market access in models where agents live and work in the same location (Redding and Sturm 2008; Bartelme 2018; Donaldson and Hornbeck 2016; Donaldson 2018;

⁵In addition, since the change in accessibility from a station depends on the geography of the city and the transit network, average treatment effects based on distance to stations in one context might not be externally valid in another. The CMA approach predicts different treatment effects from different transit networks based on the specific network structure.

Alder 2019). This paper considers a different class of urban models where individuals can live and work in separate locations. This distinction leads to meaningful differences in the way the same transit network might affect firm access to workers and resident access to jobs in any location.⁶ I use the context provided by a large, real world change in transit infrastructure to show these differential shocks to employment and residence capture the reallocation of economic activity in the city.

This paper also contributes to the growing body of work on quantitative spatial models (Ahlfeldt et. al. 2015; Allen et. al. 2015; Bird and Venables 2019; Fajgelbaum and Schaal 2020; Monte et. al. 2018; Owens et. al. 2020; Severen 2021; Bryan and Morten 2019; Heblich et. al. 2020; Adao et. al. 2019; Allen and Arkolakis 2021). Its main contribution lies in the development of a model in which multiple worker groups have non-homothetic preferences over transit modes and residential amenities. This allows the model to capture how new transit can affect the distribution of welfare across groups through their differential reliance on public transit, and through residential neighborhood choice and gentrification when house prices rise in response to better transit access.

Lastly, this paper relates to work in transportation economics measuring the benefits of improved transportation through the VTTS (Train and McFadden 1978; Small and Verhoef 2007). It connects with work measuring agglomeration externalities, providing intra-city estimates of productivity and amenity spillovers in a developing country city, identified using an expansion in the transit network that separately shifts the supply of labor and residents across the city.⁷

The paper proceeds as follows. Section 2 discusses the context of TransMilenio and the data. Section 3 develops the model. Section 4 presents and estimates the sufficient statistics approach to assess the BRT's aggregate effects. Section 5 estimates the full model to measure its distributional effects. Section 6 concludes.

2 Background and Data

2.1 TransMilenio: The World's Most-Used BRT System

Background Bogotá is the economic center of Colombia, accounting for 16% and 25% of population and GDP respectively. In 1995, the average work commute took 55 minutes, more than double that in US cities. The vast majority were taken by bus (73%), followed by car (17%) and walking (9%).⁸ Despite its importance, public transit was highly inefficient. Bus companies operated routes allocated to them by the city, but a lack of entry controls led to a large over-supply of vehicles. Low enforcement meant that up to half of the city's bus fleet operated illegally (Cracknell 2003). Disregard of bus stops led to frequent boarding and alighting along curbs, further reducing traffic flows.

⁶In trade models, firm and consumer market access often equal each other under balanced trade (e.g. Donaldson and Hornbeck 2016). One can show in my setting that it is precisely the absence of balanced trade in commuters (which would imply residence equals employment in each location, which clearly fails in the data) that delivers the BRT's very distinct shocks to resident and firm CMA shown in Figure 1.

⁷Rosenthal and Strange (2004) provide a review. Other papers using potentially exogenous sources of variation in the density of (i) employment include Combes et. al. (2010), Greenstone et. al. (2010), Kline and Moretti (2014), Ahlfeldt et. al. (2015) and (ii) residence include Bayer et. al. (2007), Guerrieri et. al. (2013), Diamond (2016), Giannone (2021).

⁸Bicycles and motorbikes account for the remaining 1% of commutes.

At the start of his first term as mayor of Bogotá, Enrique Peñalosa wasted no time in transforming the city's transit infrastructure. TransMilenio was approved in March 1998, and its first phase opened a mere 21 months later, adding 42 km along Avenida Caracas and Calle 80, two arteries of the city.⁹ Phases 2 and 3 added an additional 70km in 2006 and 2012, creating a network spanning the majority of the city. Today, the system is recognized as the “gold standard” of BRT and with more than 2.2mm riders a day using its 147 stations. It is the most heavily utilized system of its kind in the world (Cervero et. al. 2013).¹⁰ Its average operational speed of 26.2kmh reported during phase one is on par with that of the New York subway (Cracknell 2003), and far surpassed the reported 10kmh speeds on the incumbent bus network (Wright and Hook 2007).

The system involves exclusive dual bus lanes running along the median of arterial roads in the city separated from other traffic. Buses stop only at stations which are entered using a smart card so that fares are paid before arriving at platforms. Dual lanes allow for both express and local services, and passing at stations. Accessibility for poorer citizens in the urban periphery is increased through a network of feeder buses that use existing roads to bring passengers to “portals” at the end of trunk lines at no additional cost. Free transfers and a fixed fare further enhance the subsidization of the poor (at the periphery) while the government sets fares close to those offered by existing buses.

BRT is a particularly attractive alternative to subways in developing country cities since it (i) delivers similar reductions in commute times at a fraction of the cost and (ii) is much faster to build. These features have led to systems being built in more than 200 cities, the vast majority constructed over the past 15 years in Latin America and Asia (BRT Data 2017).

Route Selection and System Rollout The corridors built during the first phase of the system were consistently mentioned in 30 years of transportation studies as first-priority for mass transit (Cracknell 2003). These studies chose routes based on current and future demand levels and expected capital costs. The result was a network that connected the city center with dense residential areas in the north, northwest and south of the city (Hidalgo and Graftieux 2005). The number of car lanes was left unchanged either because existing busways were converted or due to road widening.¹¹

Three features make TransMilenio an attractive context for empirical analysis. First, since 1980 multiple administrations worked on proposals for a subway system. These can be used as placebo checks. Second, having identified neighborhoods in the city's periphery to be connected with the center, the final routes were chosen largely to minimize construction costs. Lines were placed along wide arterial roads, which were cheaper to convert and determined by the the city's historical evolution. Third, TransMilenio was rolled out quickly, primarily so that a portion could be completed within Mayor Peñalosa's term that ran between 1998 and 2001. The unanticipated nature of the system's construction and the staggered opening of lines across three phases provide sources of time

⁹While the anticipation of a system may predate its inauguration, TransMilenio went from a “general idea” to implementation in only 35 months (Hidalgo and Graftieux 2008). A “pico y placa” driving restriction implemented two years prior to TransMilenio had little impact on overall car use (Lawell et. al. 2017).

¹⁰A map of each system component and their opening date is provided in Figure A.1, while Figure A.2 shows a station before and after TransMilenio was built. For comparison, the London tube carries 5 million passengers per day over a network of 402km, giving it a daily ridership per km of 12,000 compared to TransMilenio's 20,000.

¹¹See Cracknell (2003) for discussion. This was confirmed through inspection of satellite images.

series variation used in the analysis.

One central criticism of TransMilenio was its singular focus on improving urban mobility without coordinated changes in land use regulation (Bocarejo et. al. 2013): Appendix G shows that housing supply did not respond to the system's construction. An integrated land use and transit policy that increases housing densities near stations allows more residents and firms to take advantage of improved commuting infrastructure, and sales of development rights can finance construction. In Section 5.3, I assess the impact of TransMilenio had Bogotá pursued a such an integrated policy.

Trip Characteristics Appendix G summarizes TransMilenio use. First, it is a quantitatively important mode of transit used more for longer trips than other modes. Second, TransMilenio provides an increase in door-to-door speeds of around 17% over existing buses, but remains 8.1% slower than cars. Third, when compared to other modes the BRT is used more for work commutes than leisure trips. TransMilenio's outsized role in commuting motivates the focus on access to jobs in this paper.

Yet this improvement in public transit may have differentially affected the rich and poor. Table A.19 shows that prior to TransMilenio, commutes by car were around 35% faster than bus but that low-educated Bogotanos were about 29% more likely to use buses than cars. Both facts are robust to controlling for origin-destination pair fixed effects to adjust for differences in trip composition.

2.2 Data

This section summarizes the data used in the analysis, with further details in Appendix F. The primary geographic unit used in the analysis is the census tract ("sección"). Bogotá is partitioned into 2,799 tracts, with an average size of 133,303 square meters and a mean population of 2,429 in 2005. These are contained within larger spatial units including 19 localities and 113 planning zones (UPZs).

The primary source of population data is the Department of Statistics' (DANE) General Censuses of 1993, 2005 and 2018. This provides the residential population of each block by education level. College-educated individuals are defined as those with some post-secondary education.

Employment data come from two sources. The first is a census covering the universe of establishments from DANE's 2005 General Census and 1990 Economic Census which report the location, industry and employment of each unit. The second is a database of establishments registered with the city's Chamber of Commerce (CCB) in 2000 and 2015. The data from 2015 contain the location, industry and employment of each establishment, but in 2000 employment is not provided. I therefore use establishment counts to proxy for employment, but show that establishment count and employment densities are highly correlated in years where both are available. An additional concern is that the spatial distribution of registered establishments may be different from that of total establishments. Figure A.7 shows that the employment and establishment densities in both years of the CCB data are highly correlated with the 2005 census. Coverage is even across rich and poor neighborhoods, suggesting both that the CCB data is fairly representative of overall employment. The main specifications examine changes from the CCB data, allowing employment over 10 years to respond to the first two phases of the system, but additional analyses use the economic census data to examine the

impacts of phase 1 on employment growth in the 4 years following TransMilenio's opening.

Housing market data between 2000 and 2018 come from Bogotá's Cadastre. Its mission is to keep the city's geographical information up-to-date; all parcels, formal or informal, are included and the dataset covers 98.6% of the city's more than two million properties (Ruiz and Vallejo 2010).¹² It reports the use, floorspace and land area, and value per square meter of land and floorspace, as well as a number of property characteristics. Values in the cadastre are important for the government since they determine property taxes which comprise a substantial portion of city revenue. In developed countries, these valuations are typically determined using information on market transactions. However, Bogotá, like most developing cities, lacks comprehensive records of such data and those available may be subject to systematic under-reporting. The city addresses this through an innovative approach involving sending officials to pose as potential buyers in order to negotiate a sales price under the premise of a cash payment (Anselin and Lozano-Gracia 2012). Professional assessors are also sent to value at least one property in one of each of the city's more than 16,000 "homogenous zones" (Ruiz and Vallejo 2010). As a result, Figure A.8 shows the average price per square meter of floorspace in the cadastre is highly correlated with the average purchase price per room reported in a DANE worker survey. Importantly, the relationship is constant across rich and poor neighborhoods which would not be the case were the cadastre over- or under-valuing expensive properties.

Microdata on commuting behavior come from the city's Mobility Survey administered by the Department of Mobility and overseen by DANE in 2005, 2011 and 2015. For 1995, I obtained the Mobility Survey undertaken by the Japan International Cooperation Agency (JICA) to similar specifications as the DANE surveys in later years. These are representative household surveys in which each respondent was asked to complete a travel diary for the previous day. The survey reports the demographic information of each traveler and household, including age, education, gender, industry of occupation, car ownership and in some years income. For each trip, the data report the departure time, arrival time, purpose of the trip, travel mode, and origin and destination UPZs.

Employment data by worker come from DANE's Continuing Household Survey (ECH) between 2000 and 2005, and its extension into the Integrated Household Survey (GEIH) for the 2008-2015. These are monthly, repeated cross-sectional labor market surveys covering approximately 10,000 households in Bogotá annually. Commute times between each pair of census tracts by mode are computed in ArcGIS using shapefiles of each mode's network from the city. Figure A.10 shows the computed times correlate well with observed door-to-door times from the Mobility Surveys.

3 A Quantitative Model of a City with Heterogeneous Skills

This section develops a quantitative model that captures the impact of transit infrastructure on the spatial organization of economic activity within a city. It departs from recent work (e.g. Ahlfeldt et. al. 2015) by incorporating multiple skill groups of workers, commute modes and industries.

¹²High coverage was confirmed by overlaying the shapefile for available properties over satellite images. Underlining the importance of property taxes, in 2008 they accounted for 19.8% of Bogotá's tax revenues (Uribe Sanchez 2010).

Locations differ in terms of commute times, housing floorspace, and (exogenous) amenities and productivities.¹³ Workers decide where to live, whether to own a car, where to work, and which mode of transit to use to commute. Public transit is available to everyone, but only those with a car have the option to drive. Non-homothetic preferences for car ownership and residential amenities mean the rich are more likely to own cars and live in high amenities neighborhoods. Amenities and productivities also have components that depend on local economic activity. In equilibrium, floorspace use, floorspace prices and wages adjust to clear markets.

3.1 Workers

The city is populated by worker groups indexed by $g \in G = \{L, H\}$ with a fixed population \bar{L}_g . A worker ω in group g chooses a location i in which to live, a location j in which to work, whether or not to own a car $a \in \{0, 1\}$, and the mode of transport m to use to commute to work. Individuals derive utility from consumption of a freely traded numeraire good ($C_i(\omega)$); consumption of residential floorspace ($H_{Ri}(\omega)$); and an amenity reflecting the average preference of each group to live in i under car ownership a (u_{ia}). Owning a car provides an additional mode to use for commuting and an amenity benefit, but comes at a fixed cost p_{car} . Workers are heterogeneous in their match-productivity with firms in each location ($\epsilon_j(\omega)$), their preference for each residence-car ownership pair ($\nu_{ia}(\omega)$), and their disutility from commuting that reduces their productivity at work ($d_{ijm}(\omega) \geq 1$). Land is owned by residents and rents are redistributed lump sum through payment π .¹⁴

Individuals have Stone-Geary preferences in which they need a minimum amount of floorspace \bar{h} in which to live. The indirect utility of a worker who has made choice (i, j, a, m) is then

$$U_{ijamg}(\omega) = u_{ia} \left(\frac{w_{jg} \epsilon_j(\omega)}{d_{ijm}(\omega)} - p_a a - r_{Ri} \bar{h} + \pi \right) r_{Ri}^{\beta-1} \nu_{ia}(\omega) \quad (1)$$

where w_{jg} is the wage per effective unit of labor, r_{Ri} is the price of residential floorspace in i , and $p_a = p_{car}$ for $a = 1$ and 0 otherwise.

The fixed expenditures on cars and housing allow me to match the Engel curves I document for car ownership and housing expenditure (Figure A.9) and drive sorting of workers over car ownership and residential neighborhoods by income. When cars are quicker than public transit, the rich are more willing to pay the fixed cost since their value of time is higher. The fixed expenditure on subsistence housing means that the poor spend a greater share of their income on housing and are attracted to low amenity neighborhoods where it is cheap.

Workers first choose where to live and whether or not to own a car, then where to work, and finally which transportation mode to use.¹⁵ I solve their problem by backward induction.

Mode Choice Having chosen where to live and work and whether to own a car, individuals choose

¹³Since housing supply was unaffected by TransMilenio (Appendix G), total floorspace in a location is taken as given.

¹⁴Specifically $\pi = \bar{L}^{-1} \sum_i (r_{Ri} H_{Ri} + r_{Fi} H_{Fi})$. This ensures that all the gains are accounted for within the model while avoiding inefficiencies introduced by absentee landlords that would impact the application of Proposition 2.

¹⁵While allowing for joint decisions greatly complicates inversion and estimation of the model in the presence of fixed components of expenditure and income, Appendix A.4 solves such a model and finds qualitatively similar results.

which mode of transport to use to commute to work. Commuters have nested logit demand across modes. A nest of public modes $\mathcal{B}_{Pub} \equiv \{\text{Walk, Bus, TransMilenio}\}$ is available to everyone while a nest of private modes $\mathcal{B}_{Priv} \equiv \{\text{Car}\}$ is available only to car owners. Therefore, the set of modes available depends on car ownership with $\mathcal{M}_0 = \mathcal{B}_{Pub}$ and $\mathcal{M}_1 = \mathcal{B}_{Pub} \cup \mathcal{B}_{Priv}$. Individuals have idiosyncratic preferences across modes $v_{ijm}(\omega)$ such that the realized commute cost for individual ω is given by $d_{ijm}(\omega) = \exp(\kappa t_{ijm} - b_m + v_{ijm}(\omega))$, where t_{ijm} is the time it takes to travel from i to j using mode m , b_m is a preference shifter for mode m and κ controls the mapping between commute times and costs. The commuter's problem conditional on choice (i, j, a) is simply $\min_{m \in \mathcal{M}_a} \{d_{ijm}(\omega)\}$.

Following McFadden (1974), $v_{ijm}(\omega)$ are drawn from a generalized extreme value (GEV) distribution

$$F(v_{ij1}, \dots, v_{ijN}) = 1 - \exp \left(- \sum_k \left(\sum_{m \in \mathcal{B}_k} \exp(v_{ijm}/\lambda_k) \right)^{\lambda_k} \right) \quad \text{for } k \in \{\text{Public, Private}\}.$$

This allows for correlation of preference shocks within nests, with $\lambda_k \rightarrow 0$ under perfect correlation.

Standard results for GEV distributions imply that this leads to nested logit demand for travel modes. Expected utility prior to drawing the shocks $v_{ijm}(\omega)$ is given by

$$U_{ijamg}(\omega) = u_{iag} \left(\frac{w_{jg}\epsilon_j(\omega)}{d_{ija}} - p_a a - r_{Ri} \bar{h} + \pi \right) r_{Ri}^{\beta-1} v_i(\omega)$$

where $d_{ija} = \exp(\kappa t_{ija})$ and

$$t_{ij0} = -\frac{\lambda}{\kappa} \ln \sum_{m \in \mathcal{B}_{Public}} \exp \left(b_m - \frac{\kappa}{\lambda} t_{ijm} \right) \quad (2)$$

$$t_{ij1} = -\frac{1}{\kappa} \ln (\exp(b_{car} - \kappa t_{ijCar}) + \exp(\kappa t_{ij0})). \quad (3)$$

Intuitively, the expected commute cost d_{ija} can be expressed as the inclusive value of commute times available to the individual with car ownership a .¹⁶

Employment Choice Having chosen where to live and whether to own a car, individuals draw a vector of match-productivities with firms across the city iid from a Frechet distribution $F(\epsilon_j) = \exp(-\tilde{T}_g \epsilon_j^{-\theta_g})$. Here θ_g measures the dispersion of productivities while \tilde{T}_g controls their level.

With these draws in hand, linearity of (1) means that workers choose to work in the location that offers the highest income net of commute costs $\max_j \{w_{jg}\epsilon_j(\omega)/d_{ija}\}$. Standard results imply that the probability a type- g worker who has made choice (i, a) decides to work in j is given by

$$\pi_{j|iag} = \frac{(w_{jg}/d_{ija})^{\theta_g}}{\sum_s (w_{sg}/d_{isa})^{\theta_g}} \equiv \frac{(w_{jg}/d_{ija})^{\theta_g}}{\Phi_{Riag}}. \quad (4)$$

¹⁶The set of shifters are normalized so that $\bar{t}_{iia} = 0 \forall i, a$. This is equivalent to always taking a weighted average travel time of available modes, where the weights are the preference shifters b_m .

The term $\Phi_{Riag} \equiv \sum_s (w_{sg}/d_{isa})^{\theta_g}$ is defined as Residential Commuter Market Access (RCMA). It captures residents' access to well-paid jobs from location i . Individuals are more likely to commute to locations with a high wage net of commute costs (the numerator) relative to other locations (the denominator). The sensitivity of commute decisions to commute costs is governed by the dispersion of productivities, with a greater dispersion (lower θ_g) making choices less sensitive. Differences in θ_g across groups will be important in determining the incidence of improved infrastructure, since it controls the extent to which individuals are willing to bear high commute costs to work in a location.

Expected income prior to drawing the vector of match productivities is related to RCMA through

$$\bar{y}_{iag} = T_g \Phi_{Riag}^{1/\theta_g}, \quad (5)$$

where T_g is a transformation of the location parameter of the Frechet distribution.¹⁷

Residence and Car Ownership Choice In the first stage, individuals choose where to live and whether or not to own a car to maximize expected indirect utility. The idiosyncratic preferences $\nu_{ia}(\omega)$ are drawn from a Frechet distribution with shape parameter $\eta_g > 1$ and unit scale. The supply of type- g individuals to location i and car ownership a is then

$$L_{Riag} = \lambda_{U,g} \left(u_{iag} \tilde{y}_{iag} r_{Ri}^{\beta-1} \right)^{\eta_g} \quad (6)$$

where $\tilde{y}_{iag} \equiv \bar{y}_{iag} - p_a a - r_{Ri} \bar{h} + \pi$ is expected net income and $\lambda_{U,g}$ is an equilibrium constant.¹⁸

3.1.1 Aggregation

Firm Commuter Market Access and Labor Supply The supply of workers to any location is found by summing over the number of residents who commute there $L_{Fjg} = \sum_{i,a} \pi_{j|iag} L_{Riag}$. This implies

$$L_{Fjg} = w_{jg}^{\theta_g} \Phi_{Fjg} \quad (7)$$

where $\Phi_{Fjg} = \sum_{i,a} d_{ija}^{-\theta_g} \frac{L_{Riag}}{\Phi_{Riag}}$

Labor supply is log-linear and depends on two forces. First, more workers commute to destinations paying higher wages. Second, conditional on wages firms attract workers when they have better access to them through the commuting network. This is captured through Φ_{Fjg} which I refer to as Firm Commuter Market Access as it reflects firms' access to workers. This is because individuals care about wages net of commute costs. Total effective labor supply to a location is given by

¹⁷The constants in this section are $T_g \equiv \gamma_{\theta,g} \bar{I}_g^{1/\theta_g}$, $\gamma_{\theta,g} = \Gamma\left(1 - \frac{1}{\theta_g}\right)$, $\lambda_{U,g} = \bar{L}_g (\gamma_{\eta,g} / \bar{U}_g)^{\eta_g}$ and $\gamma_{\eta,g} = \Gamma\left(1 - \frac{1}{\eta_g}\right)$ where $\Gamma(\cdot)$ is the gamma function. Expected utility prior to learning match productivities is $U_{iag}(\omega) = u_{iag} (\bar{y}_{iag} - p_a a - r_{Ri} \bar{h}) r_{Ri}^{\beta-1} \nu_{ia}(\omega)$.

¹⁸The model requires that $\pi > p_a a + r_{Ri} \bar{h} \forall i$ such that $\bar{u}_{iag} > 0$, since the Frechet distribution implies there will always be a positive mass of individuals with income arbitrarily close to zero. This is satisfied when the model is taken to the data.

$\tilde{L}_{Fjg} = \bar{\epsilon}_{jg} L_{Fjg}$, where $\bar{\epsilon}_{jg}$ is the average productivity of type- g workers who decide to work in j .¹⁹

Worker Welfare Properties of the Frechet distribution imply that average welfare in each location is equal to the expected utility prior to the first stage given by

$$\bar{U}_g = \gamma_{\eta,g} \left[\sum_{i,a} \left(u_{iag} \tilde{y}_{iag} r_{Ri}^{\beta-1} \right)^{\eta_g} \right]^{1/\eta_g} \quad (8)$$

3.1.2 Firms

Technology There are $s \in \{1, \dots, S\}$ industries that produce varieties differentiated by location under perfect competition. Output is freely traded, and consumers aggregate each variety into the numeraire under CES with elasticity of substitution $\sigma_D > 1$. Firms produce using a Cobb-Douglas technology over labor and commercial floorspace

$$Y_{js} = A_{js} N_{js}^{\alpha_s} H_{Fjs}^{1-\alpha_s}$$

where $N_{js} = \left(\sum_g \alpha_{sg} \tilde{L}_{Fjgs}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$.

The labor input is a CES aggregate over each skill group's effective labor with elasticity of substitution σ , $\alpha_s = \sum_g \alpha_{sg}$ is the total labor share, and A_{js} is the productivity of location j for firms in industry s which firms take as given.

Industries differ in terms of the intensity in which they use different types of workers α_{sg} . All else equal, industries like real estate and financial services demand more high-skilled workers while others, such as hotels and restaurants, rely on the low-skilled.

Factor Demand Perfect competition implies that the price of each variety is equal to its marginal cost $p_{js} = W_{js}^{\alpha_s} r_{Fj}^{1-\alpha_s} / A_{js}$, where r_{Fj} is the price of commercial floorspace in j and

$$W_{js} = \left(\sum_g \alpha_{sg} w_{jg}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$$

is the cost of labor for firms of industry s in location j . Intuitively, labor costs differ by industry due to their differential skill requirements. Solving the firm's cost minimization problem and letting X_{js} denote firm sales, the demand for labor and commercial floorspace is²⁰

$$\tilde{L}_{Fjgs} = \left(\frac{w_{jg}}{\alpha_{sg} W_{js}} \right)^{-\sigma} N_{js} \quad (9)$$

$$H_{Fjs} = (1 - \alpha_s) \frac{X_{js}}{r_{Fj}}. \quad (10)$$

¹⁹In particular, $\bar{\epsilon}_{jg} = T_g \sum_{i,a} \frac{\pi_j^{-1/\theta_g} \pi_{j|iag} L_{Riag}}{d_{ija} \sum_{r,o} \pi_{j|rog} L_{Rrog}}$.

²⁰Here $X_{js} = p_{js}^{1-\sigma_D} X$ where $X = \sum_i \beta (E_i - \bar{h} r_{Ri} L_{Ri})$ is total spending on goods in the city and $E_i = \sum_{g,a} (\bar{y}_{iag} - p_a a + \pi) L_{Riag}$ is total spending on goods and housing from residents in i .

3.1.3 Floorspace

Market Clearing In each location there is a fixed amount of floorspace H_i , a fraction ϑ_i of which is allocated to residential use and $1 - \vartheta_i$ to commercial use. Market clearing for residential floorspace requires its supply $H_{Ri} = \vartheta_i H_i$ equals demand:

$$r_{Ri} = (1 - \beta) \frac{E_i}{H_{Ri} - \beta \bar{h} L_{Ri}} \quad (11)$$

where $L_{Ri} = \sum_{g,a} L_{Riag}$ is the total number of residents in i . Likewise, the supply of commercial floorspace $H_{Fj} = (1 - \vartheta_i) H_j$ must equal that which is demanded by firms:

$$r_{Fj} = \frac{\sum_s (1 - \alpha_s) \left(W_{js}^{\alpha_s} r_{Fj}^{1-\alpha_s} / A_{js} \right)^{1-\varsigma} X}{H_{Fj}}. \quad (12)$$

Floorspace Use Allocation Landowners allocate floorspace ϑ_i to maximize profits. They receive r_{Ri} per unit allocated to residential use, but land use regulations limit the return to each unit allocated to commercial use to $(1 - \tau_i) r_{Fi}$. Since they maximize profits, we have

$$\begin{aligned} \vartheta_i &= 1 \text{ if } r_{Ri} > (1 - \tau_i) r_{Fi} \\ (1 - \tau_i) r_{Fi} &= r_{Ri} \quad \forall \{i : \vartheta_i \in (0, 1)\} \\ \vartheta_i &= 0 \text{ if } (1 - \tau_i) r_{Fi} > r_{Ri} \end{aligned} \quad (13)$$

3.1.4 Externalities

Productivities A location's productivity depends on both an exogenous component, \bar{A}_{js} , which reflects features independent of economic activity (e.g. access to roads, slope of land), and the density of employment in that location

$$A_{js} = \bar{A}_{js} \left(\tilde{L}_{Fj} / T_j \right)^{\mu_A}, \quad (14)$$

where $\tilde{L}_{Fj} = \sum_s \tilde{L}_{Fjs}$ is the total effective labor supplied to that location and T_j is its land area. The strength of agglomeration externalities is governed by the parameter μ_A .

Amenities Amenities depend on an exogenous component, \bar{u}_{iag} , which varies by car ownership (e.g. leafy streets, proximity to getaways surrounding the city) and a residential externality that depends on the college share of residents

$$u_{iag} = \bar{u}_{iag} \left(L_{RiH} / L_{Ri} \right)^{\mu_{U,g}}. \quad (15)$$

In contrast to existing urban models (e.g. Ahlfeldt et. al. 2015), endogenous amenities depend on demographic composition across skill groups rather than the total density of residents. This seems especially applicable in developing country cities that lack strong public goods provision. In Bogotá, where crime is a significant problem, the rich often pay for private security around their buildings

which increases the sense of safety in those areas. This externality provides an additional force toward residential segregation, since the high-skilled are more willing to pay to live in high-amenity neighborhoods and by doing so increase amenities further.²¹

3.1.5 Equilibrium

Definition. Given vectors of exogenous location characteristics $\{H_i, \bar{u}_{iag}, \bar{A}_{js}, t_{ija}, \tau_i\}$, city group-wise populations $\{\bar{L}_g\}$ and model parameters $\{\bar{h}, \beta, \alpha, p_a, \kappa, \theta_g, T_g, \eta_g, \alpha_{sg}, \sigma_D, \sigma, \mu_A, \mu_U\}$, an equilibrium is defined as a vector of endogenous objects $\{L_{Riag}, L_{Fjg}, w_{jg}, r_{Ri}, r_{Fi}, \vartheta_i, \bar{U}_g, \pi\}$ such that

1. **Labor Market Clearing** The supply of labor by individuals (7) is consistent with demand for labor by firms (9);
2. **Floorspace Market Clearing** The market for residential floorspace clears (11) and its price is consistent with residential populations (6), the market for commercial floorspace clears (12) and floorspace shares are consistent with land owner optimality (13);
3. **Closed City Populations** add up to the city total, i.e. $\bar{L}_g = \sum_{i,a} L_{Riag} \forall g$, and rents are redistributed lump sum to residents.

4 Empirical Analysis and Aggregate Effects

This section turns to a reduced-form analysis of how TransMilenio reshaped the organization of economic activity in Bogotá. To guide this analysis, I use the insight that a special case of the model delivers a log-linear reduced form between CMA and endogenous variables. The coefficients of these regressions are in fact sufficient statistics to analyze the impact of transit on the distribution of economic activity across the city and aggregate welfare. This approach has appeal in that it can speak to the BRT's aggregate effects, as its parameters can be transparently estimated via reduced form and because it holds for a broad class of models. I then turn to estimating the full model in Section 5 to measure its distributional impacts, which is of specific interest in this paper.

4.1 Reduced Form in a Special Case of the Model

Proposition 1 in Appendix C.2 shows that when there is one group of workers and firms and no fixed elements of expenditure or income, the equilibrium can be written as

$$\ln \hat{y}_{Ri} = \beta_R \ln \hat{\Phi}_{Ri} + e_{Ri} \quad (16)$$

$$\ln \hat{y}_{Fi} = \beta_F \ln \hat{\Phi}_{Fi} + e_{Fi}. \quad (17)$$

The outcome variables $\hat{y}_{Ri} = [\hat{L}_{Ri}, \hat{r}_{Ri}]$ and $\hat{y}_{Fi} = [\hat{L}_{Fi}, \hat{r}_{Fi}]$ are changes in residential and commercial

²¹TransMilenio may directly impact productivities \bar{A}_{js} and amenities \bar{u}_{iag} (e.g. through street improvements, crime, pollution). My empirical results which control for distance to station (which captures such effects) show little effect on CMA coefficients, motivating their exclusion from the model. Since Ahlfeldt et. al. (2015) find amenity and productivity spillovers decay rapidly across a few blocks, and the median tract contains 14 blocks, I omit cross-tract effects.

outcomes, consisting of residential population \hat{L}_{Ri} , residential floorspace prices \hat{r}_{Ri} , employment \hat{L}_{Fi} and commercial floorspace prices \hat{r}_{Fi} . The right-hand side variables $\hat{\Phi}_{Ri}$ and $\hat{\Phi}_{Fi}$ are changes in CMA. The coefficients $\beta_R = [\beta_{LR}, \beta_{rR}]$ and $\beta_F = [\beta_{LF}, \beta_{rF}]$ reflect both the direct and indirect effects of improving CMA as it filters through land and labor markets. Finally, the residuals contain clusters of unobserved location characteristics that are exogenous to economic activity. For residential outcomes, e_{Ri} contains changes in amenities and residential floorspace supplies while for commercial outcomes, e_{Fi} contains changes in productivities and commercial floorspace supplies.²²

This system shows that the transit network only matters for equilibrium outcomes through the two CMA variables. In fact, the change in the entire distribution of economic activity across the city only depends on the change in CMA and on a structural residual that reflects changing exogenous location fundamentals (productivities, amenities and floorspace supplies).²³ Proposition 1 also shows that this specification is shared by a broad class of urban models that include housing supply, firm mobility, capital as a productive input, and leisure in utility. Moreover, it shows that (i) the change in CMA terms and (ii) the reduced-form elasticities of outcomes to CMA are sufficient statistics for the *relative* change in economic activity across the city in response to changes in transit infrastructure. The overall *level* of the changes are pinned down by an assumption on population mobility into the city from the rest of the country, as well as by values for two parameters σ and β that cannot be estimated from the reduced form. These must be specified in some other way by the researcher, for example by calibrating them to external values or aggregate moments.

The CMA terms can be easily recovered (to scale) using data on residential populations, employment, commute costs d_{ij} , and the commuting elasticity θ from the following system of equations²⁴

$$\Phi_{Ri} = \sum_j d_{ij}^{-\theta} \frac{L_{Fj}}{\Phi_{Fj}} \quad (18)$$

$$\Phi_{Fj} = \sum_i d_{ij}^{-\theta} \frac{L_{Ri}}{\Phi_{Ri}}. \quad (19)$$

RCMA reflects access to well-paid jobs. It is greater when a location is close (in terms of having low commute costs) to other locations with high employment, particularly so when these other locations lack access to workers (increasing the wages that firms there are willing to pay). FCMA reflects access to workers through the commuting network. It is greater when a location is close to other locations with high residential populations, particularly so when these other locations lack access to jobs (lowering the wages that individuals are willing to work for there).

²²As discussed in Appendix C.2, there is one first order approximation $\tilde{\Phi}_{Fi} \approx \Phi_{Fi}^{\frac{\theta-1}{\theta}}$ involved to get this reduced form. The reduced form this approximation comes from involves $\Phi_{Ri}, \Phi_{Fi}, \tilde{\Phi}_{Fi}$ on the right hand side and holds exactly. Since $\Phi_{Fi}, \tilde{\Phi}_{Fi}$ are very highly correlated in the data (correlation coefficient of 0.98), this is not empirically feasible to implement and so the approximation is used to generate this simpler reduced form. The coefficients can be mapped between the two, however, so that the unapproximated version is used to conduct counterfactuals.

²³The contents of the residual and reduced form parameters are outlined in Appendix C.7. All contents in the residual are exogenous to the model, except two equilibrium constants (average utility and total expenditure) that are absorbed into the regression constant.

²⁴In this single-group model, there is a single commute elasticity θ .

I now turn to the specific context of Bogotá to visualize the change in CMA and how it differs from the distance-based measures of treatment effects commonly found in the literature. Figure 1 plots the distribution of changes in CMA induced by the construction of the first two phases of the system.²⁵ The system increased access to jobs much more in the outskirts of the city, which were far from the high-employment densities in the center. Firms’ access to workers rose more in the center, where firms stood to benefit from the increased labor supply along all spokes of the network.

4.2 Measuring CMA

Computing changes in CMA induced by TransMilenio requires values for θ and $d_{ij} = \exp(\kappa t_{ij})$.

Identifying κ, λ, b_m . The mode choice parameters are estimated via maximum likelihood using standard expressions for choice shares in the nested logit model from Section 3.1 (see Appendix C.3). The data come from the 2015 Mobility Survey when all modes are available. κ is identified from the sensitivity of choices to differences in travel time across options, λ is identified from the differential sensitivity within public modes, and the preference shifters b_m are identified from differences in choice shares conditional on observed travel times. The results are in Panel A of Table 1. The estimate of $\kappa = 0.011$ is very close to the 0.01 reported in Ahlfeldt et. al. (2015). The value $\lambda = 0.157$ indicates a sizable correlation of draws within the public nest. Conditional on travel time, cars are most attractive, followed by buses and TransMilenio. That TransMilenio is least desirable likely reflects high crowds on the system, and the inconvenience of having to walk between stations and final origins and destinations. With these parameters, t_{ij0} and t_{ij1} can be obtained from (2) and (3). However, in the simple model considered in this section there is no car ownership. As described in Appendix C.4, average commute times t_{ij} are computed by assuming residents become car owners according to a Bernoulli distribution, with probability equal to the share of car owners in Bogotá.

Identifying θ . As shown in Appendix C.4, the special case of the model yields a simple gravity equation that identifies the parameter cluster $\theta\kappa$ from the sensitivity of the change in commute flows to the change in commute times between any two locations, controlling for origin and destination fixed effects. Estimating this relationship via PPML to account for zeros in the data yields a value of $\theta = 3.398$ in Panel B of Table 1, similar to existing estimates (Monte et. al. 2018; Heblich et. al. 2020). I use this as the baseline value, but use alternative values in robustness checks (see Appendix C.4).

4.3 Empirical Results

Identifying TransMilenio’s Reduced Form Effect on Economic Activity. The regressions (16) and (17) may be biased if Bogotá’s government chose routes in a way that targeted neighborhoods with differential trends in unobserved characteristics (such as if trying to stimulate lagging regions or to

²⁵The figure plots the change in CMA holding population and employment fixed at their initial level in 1993 and 1990 respectively and changing only commute costs. This isolates the change due only to TransMilenio (discussed in Section 4.3). FCMA increases toward the center-North due to the high density of (low-skilled) workers in the South.

support thriving ones). Instead of leaning on a single approach, I exploit a variety of TransMilenio’s institutional features to establish its causal impact on Bogotá’s structure.

First, I include a rich set of controls, including locality fixed effects to (partially) control for changes in unobservables. Second, I use variation in CMA induced by changes in the network more than 1.5km from a location, which is less likely to be correlated with changes in surrounding unobservables. Third, I condition on distance to closest TransMilenio station to assess whether the effects are driven by changes in accessibility rather than by other station features (e.g. changes in foot traffic, pollution or complementary infrastructure). Fourth, I digitize four different historical plans for Bogotá’s transit network and run specifications including both the realized change in CMA and the change induced by these (hypothetical) planned networks. The coefficients on the planned CMA variables can be interpreted as a placebo check that the planned-but-unbuilt locations do not grow differentially in the absence of new transit. The stability of the coefficients on the realized CMA variables addresses any omitted variable bias (not captured by the controls) that can arise from a location’s non-random exposure to transport infrastructure (Borusyak and Hull 2021). Fifth, I exploit TransMilenio’s staggered rollout across three phases by using event studies and falsification tests which assess whether there is growth in outcomes prior to line openings. Sixth, I construct cost-shifting instruments to predict TransMilenio’s routes based on a historical tram network and on engineering estimates of the cost to build BRT on different types of land.

An additional challenge is that changes in CMA contain population and employment in both periods. Since productivity and amenity shocks that determine residential population and employment are contained in the error terms, they will be mechanically correlated with changes in CMA. I thus construct versions of the change in CMA by solving (18) and (19) while holding population and employment fixed at their initial levels, allowing only commute costs to change, and use these throughout the empirical analysis. This isolates the variation in CMA due only to changing commute costs through TransMilenio’s construction. After solving for the CMA terms, I construct the change in CMA for a given location by excluding the location itself in the summation. This addresses the possibility that changes in unobservables may be correlated with a location’s initial level of economic activity. The main specifications use these CMA measures as regressors, but later in this section I use these to instrument for the “realized” change in CMA.

Main Specification. Table 2 presents the baseline results. Each entry corresponds to the coefficient from a regression of the change in each outcome on the change in RCMA or FCMA in each census tract. Since the data do not all line up, each specification relies on changes over different periods. However, the changes in CMA are always measured using changes in commute times due to TransMilenio routes constructed between the two periods over which the outcome is measured.²⁶ Es-

²⁶In population regressions, the outcome is the log change in residential population between 1993 and 2018. The change in CMA is that induced by all three phases of TransMilenio, holding residential population and employment fixed at their levels in 1993 and 1990 respectively. In land market regressions, outcomes are log changes between 2018 and 2000 and the change in CMA is that induced by all three phases holding residential population and employment fixed at their levels in 2000 (population in 2000 is a linear interpolation from the 1993 and 2005 census; employment is from the 2000 CCB data). Establishment regressions regress changes between 2000 and 2015 from the CCB data against the same CMA measures

establishment regressions are weighted by the share of establishments in 2000 in each tract to increase precision.²⁷ Since some establishment results are noisy, I include the share of floorspace used for commercial purposes as an outcome to provide supplemental evidence for TransMilenio’s impact on the reallocation of employment.²⁸

Column (1) includes controls for locality fixed effects, basic tract characteristics, and log distance to CBD interacted with region dummies.²⁹ Changes in CMA due to TransMilenio have strong, positive impacts on all outcomes. These relationships remain mostly stable as more controls are added in columns (2) and (3), sometimes becoming sharper. The exception is log establishments in the final row, whose coefficient falls by a third with the full set of controls. I consider column (3) to be the baseline specification continued in later tables, as it includes the full set of controls.

Column (4) excludes tracts that are closer than 500m from an endpoint of a TransMilenio route (a “portal”) or the CBD. The intent of the government was to connect outlying neighborhoods with the CBD, so the location of these portals may have been endogenous to underlying trends in local economic activity. The coefficients remain largely stable in this subsample of tracts, suggesting that endogeneity in the locations directly targeted by TransMilenio is not driving the results.

Column (5) uses the change in CMA to locations farther than 1.5km away from a tract. Network additions at this distance from a tract are less likely to be linked to local trends in unobservables. The results remain robust and, for the most part, stable. Column (6) assumes users take the quickest mode of public transit available, and shows the results are robust to alternative forms of aggregation.

Lastly, column (7) conditions on distance to stations to establish that the effects are primarily driven by changes in accessibility rather than by station features (e.g. changes in foot traffic, pollution or complementary infrastructure). This finding supports the model’s emphasis on accessibility.

Visualizing the Relationship. Figure 2 plots the non-parametric relationship between residual growth in outcomes and CMA. The relationship appears approximately log-linear for each outcome, as predicted by the model. The simple model seems to capture the heterogeneous effects observed in the data: tracts with large improvements in accessibility experience large growth in outcomes.

Hypothetical Changes in CMA from Historical Network Plans. The location of the TransMilenio network was not random. The government may have located the network to support or spur existing local trends in economic activity. To provide additional evidence of TransMilenio’s causal impact, I leverage four distinct historical plans for a transit network digitized from planning documents.

as the land market regressions. This is preferred to the census employment data since it gives employment 10 years to respond to TransMilenio. Table 4 uses employment data from the census to examine the impact of TransMilenio lines built during phase 1 (by 2003) on employment growth between 1990 and 2005.

²⁷Unweighted regressions are presented in Table A.5.

²⁸In the simple model this should not change. An extension in Appendix E.3 allows for floorspace use shares to respond to TransMilenio (with total floorspace supply held constant, as observed in the data), which delivers heterogeneous elasticities of economic activity to CMA. The model with endogenous housing supply in Appendix C.6 also allows for endogenous floorspace use shares (via changes in relative supplies for residential and commercial floorspace) and admits the same reduced form as the baseline model.

²⁹The North is richer and more educated than the West and South, so this allows for differential growth further away from the city center within each region.

Since 1980, multiple administrations had worked on proposals for building a subway or metro system in Bogotá. Four distinct plans for the network were prepared before Mayor Peñalosa agreed with the proposal by JICA to build a BRT, given that the cost of a subway would have been “ten times higher than the alternative of articulated buses”. I obtained and digitized the maps for these four planned networks, shown in Figure A.3.³⁰ I then solve for the predicted change in CMA had TransMilenio been built along each of the planned networks, and compute the average change in log RCMA and FCMA across all four plans.

The baseline specification (column 3 of Table 2) is then extended to include these expected changes in CMA under the plans as additional regressors. One interpretation of the results is as a placebo check. If the observed impacts are due to TransMilenio itself rather than to the selection of routes based on trends in unobservables in adjacent neighborhoods, there should be no impact of these planned-but-unbuilt networks. A second interpretation is that this controls for the omitted variable bias that can arise from a location’s non-random exposure to transport infrastructure, as highlighted by Borusyak and Hull (2021). The idea is that some locations may receive systemically different changes in accessibility under any network realization. For example, central neighborhoods will tend to have greater increases in FCMA since they are close to where workers live by virtue of their central location. Identification requires that these “on average” more exposed locations do not differ in their trends in unobservables. While any such trends may already be controlled for by the rich fixed effects and controls used in this paper, controlling for the average change in CMA under these counterfactual networks conducts the exact “recentering” shown by the authors to remove the omitted variable bias. If the controls already capture any differential trend in unobservables in more “on average” exposed locations, then the coefficient on the realized CMA terms should be invariant to the inclusion of the expected change in CMA and the coefficient on the latter should be zero.

Table 3 presents the results. Column (1) repeats the baseline specification, while column (2) adds the control for the expected change in CMA across the four plans. In each case, the coefficient on the realized change in CMA due to TransMilenio is invariant to the inclusion of this additional control. The p-value testing for equality of coefficients on the realized CMA variable across both columns ranges from 0.24 to 0.96. The coefficient on the planned CMA variable is statistically insignificant from zero in all specifications. These results suggest two things: first, that the observed impacts of TransMilenio are unlikely due to pre-existing trends in neighborhoods selected by city planners, and second, that the existing set of controls does a sufficient job in controlling for any omitted variable bias that could arise from non-random exposure to the network.

Staggered Station Openings. TransMilenio was opened in three distinct phases during the 2000s and 2010s.³¹ This section runs a set of falsification checks to test for changes in outcomes prior to the

³⁰See (“Historia de TransMilenio”) for the quote, and Alcaldía Mayor de Bogotá D.C. (2009) for the network maps. As discussed in Appendix E.3, I add predicted feeder routes under these networks by placing a 2km radius disk around each end point of the planned lines connecting the two with 8 “spokes”, and create stops every 250m.

³¹Phase 1 consisted of 3 lines in 2000, and 1 line each in 2001, 2002 and 2003. Each year consisted of 47%, 26%, 6% and 21% of the stations opened in phase 1, respectively. Phase 2 consisted of 2 lines in 2005 and 1 line in 2006, with each year accounting for roughly half the stations opened in this phase. Lastly, phase 3 consisted of 2 lines in 2012 and 1 in 2013 with

opening of stations in later phases. The specification is

$$\Delta_{t,t-\ell} \ln y_i = \beta^{Current} \Delta_{t,t-\ell} \ln \Phi_i + \beta^{Future} \Delta_{t+k,t} \ln \Phi_i + \gamma' X_i + \varepsilon_i$$

The outcome is the growth in a variable, y_i , between two periods, t and $t - \ell$ (e.g. 2006 and 2000). This is regressed on (i) the change in CMA between t and $t - \ell$, (ii) the change in future CMA between $t + k$ and t (e.g. 2015 and 2006), as well as the same set of controls as the baseline specification. If there is no growth in outcomes prior to TransMilenio being built, the coefficient β^{Future} should be zero.

The time periods are chosen to best line up with the available data and the opening of TransMilenio lines. Since the openings of phases 1 and 2 are spread out between 2000 and 2006 (with every year except 2004 experiencing station openings), I focus the analysis on phase 3, which opened in 2012 and 2013. For land market outcomes, the change in outcomes is measured between 2008 and 2000. The right-hand side variables include CMA growth due to (i) phases 1 and 2 of the system open by 2006 (to identify $\beta^{Current}$) and (ii) phase 3 of the system open by 2013 (to identify β^{Future}). While prices may experience some anticipation effects, plans for phase 3 were mired by uncertainty and delays, with construction only beginning in late 2009. For residential population, the change is measured between the 1993 and 2005 censuses. The right-hand side variables include CMA growth due to phase 1 (open by 2003, with most stations opening by 2001), as well as the change in CMA due to phases 2 and 3. Lastly, for employment, I turn to the measures from the economic census rather than the CCB data. While the latter is available only in 2000 and 2015 (bookending the entire network construction), the economic census is available in 1990 and 2005. This permits me to separately examine the impacts of changes in CMA due to phase 1 versus phases 2 and 3, similar to residential population.³²

In Table 4, Panel A presents the results for residential population and residential floorspace prices. Odd columns repeat the baseline specification but with outcomes measured over this different period (e.g. 1993 to 2005 for residential population, compared to 1993 to 2018 in the baseline results). The positive relationships remain significant, although the point estimates are somewhat attenuated. This might be expected given that there is less time for outcomes to respond to the change in CMA than in the baseline specification. Even columns then run the specification above. They maintain a significant relationship between outcome growth and CMA growth due to lines constructed over the period, but an insignificant impact due to accessibility from future lines. While insignificant, these estimates of β^{Future} can be noisy. Panel B finds similar patterns for commercial land market outcomes.

Panel C presents the impact on total and formal employment from the economic census.³³ In the odd columns, which regress on realized changes in CMA, I document positive but noisy coefficients (p-values of 0.123 and 0.118). These estimates are larger than the baseline estimates using the CCB data, but the difference is statistically insignificant given the imprecision of the estimates. The even

2012 representing 84% of the openings

³²The downside of the economic census data is that there is less time for employment to adjust: on average across lines opening during phase 1, there are just under 4 years between the opening year and the 2005 economic census. This compares to 10.5 years between the average opening year and the 2015 CCB.

³³Formal employment is defined as employment in establishments with 5 or more workers.

columns add in future CMA growth, which is statistically insignificant in both cases.

Floorspace Price Event Study. I leverage the annual cadastral data to examine more granular house price dynamics prior to the opening of TransMilenio’s third phase. I run regressions of the form

$$\ln r_{Rit} = \alpha_i + \gamma_{\ell(i)t} + \sum_{\tau=-8}^{\tau=6} \beta^\tau \Delta_{12,06} \ln \Phi_{Ri} + \delta'_t X_i + \varepsilon_{it},$$

where α_i are tract fixed effects, $\gamma_{\ell(i)t}$ are locality-year fixed effects, and $\delta'_t X_i$ is a set of controls with time-varying coefficients. The controls include those from the baseline specification, but add the change in CMA due to the first two phases of the system to capture the impact of changes that these earlier lines had on house prices that is correlated with the change due to phase 3. The regression is weighted by initial floorspace price in 2000 to improve precision. The β^τ coefficients capture the response of residential floorspace prices in a tract τ years from the third phase lines opening to the change in CMA due to the lines that open during this phase.

Figure 3 plots the event study coefficients. Reassuringly, the change in CMA induced by the network expansion in phase 3 has no impact on floorspace price growth to the line openings. It is not clear ex ante that this would be the case. Prices could rise due to anticipatory effects as expectations around whether and where the line would open firm up. Alternatively they could fall due to the disamenities surrounding the construction from late 2009 through 2012. In fact, consistent with this possibility, there is a mild decrease in house prices in tracts that experienced a larger growth in accessibility due to phase 3 in the two or three years prior to opening. The year before the lines open, the responsiveness of prices to CMA jumps approximately 0.4 log points. This is potentially due to anticipation effects as the opening of the third phase became certain.³⁴ This effect is stable until two years after opening, after which the elasticity rises 1 log point until six years after opening.

While the difference between the short- and medium-run effects in this event study may reflect the multiplier effect due to the reallocation of population and employment to treated areas, it is important to note that (part of) this could also be due to the way the data is constructed. As described in Appendix F, part of the annual change in prices in the cadastral database is based on inflating prior years’ values. Primary data are collected by the cadastral office to fully reassess properties—based on collecting information on properties for sale, making offers to elicit true sales values, and having in person visits by professional assessors—but this happens fairly infrequently (around three times over the period in question). This motivates the focus on long-run impacts in the rest of the paper.

Instrumental Variables to Predict TransMilenio’s Placement. Lastly, I construct two cost-shifting instruments for TransMilenio routes. These in turn imply two instruments for the change in CMA.³⁵

³⁴Corruption cases surrounding the construction of the third phase had added to construction delays, which may have brought more uncertainty than usual to whether and when the lines would actually open.

³⁵Additional details can be found in Appendix F.3. To compute the instruments, I first calculate the commute times had the system been built along each instrument. Plugging these into (18) and (19) and continuing to hold population and employment fixed at their initial level, I obtain the predicted CMA had TransMilenio been built along these routes. My instrument for the change in CMA is then the difference between this predicted CMA under TransMilenio and its value

The first takes as given the government’s overall strategy of connecting portals at the edge of the city with the CBD, excludes those areas from the analysis, and constructs the routes that would have been built if the sole aim had been to minimize costs. This is done by using engineering estimates to compute the cost to build BRT in each parcel of land in Bogotá based on its land use in 1980. This is a valid instrument when these least-cost routes predict TransMilenio’s placement but are uncorrelated with trends in unobserved amenities and productivities (conditional on controls).

The second instrument exploits the location of a tram system that opened in 1884, which was last extended in 1921 and stopped operating in 1951. I extend the 1921 lines to the present edge of the city to improve predictive fit, given the city’s substantial expansion over the period. The tram was built along wide arterial roads, which are cheaper to convert to BRT than narrow ones. The tram may have had persistent direct effects on trends in unobservables that lasted well after its construction, which I capture by including historical controls. Conditional on these historical variables, the tram routes should be uncorrelated with changes in productivities and amenities between 2000 and 2012 to the extent that these were unanticipated by city planners in 1921.

The identification assumption is that the instruments have only an indirect effect on outcome growth through the predicted change in CMA. One concern is that features that make a location cheaper to build BRT, such as proximity to a main road, can have direct effects on outcomes. A key advantage of my approach is that I can control for distance to these features (distance to the tram, distance to main roads) and use only residual variation in predicted CMA growth for identification.

Table 5 presents the results. Column 1 reproduces the baseline results for reference. Column 2 shows the results are very similar when instrumenting the realized change in CMA (allowing residence and employment to change across periods and summing over all locations) with the measure from the baseline specification. Columns 3 and 4 instrument the realized change in CMA using the average change across the tram and least-cost path instruments, using either all tracts except the tract itself (column 3) or only tracts 1.5km away (column 4). The coefficients are mostly stable across these specifications, with the exception of residential floorspace prices which roughly double when moving to the final column. The population and commercial floorspace price coefficients are imprecise, but sharpen in the last column. Taken together, these specifications support the impression from the analyses above—that changes in CMA due to TransMilenio seem unrelated to trends in unobservables conditional on the rich set of controls. Given the broad stability of the estimates across specifications, I use the coefficients from the baseline specification in the next section but explore the robustness of the results to using the elasticities from the other columns in this table (see Table A.2).

Robustness Checks and Additional Results. Appendix G.6 presents robustness of these results to alternative i) methods of aggregating times; ii) commute elasticities; iii) clustering of standard errors; iv) additional controls and sample selection criterion and v) weighting procedures. Appendix G.4 provides evidence that TransMilenio increased wages but also led to a sorting response where the high-skilled moved into neighborhoods with improved market access. This is consistent with the

in the initial period without the system. Historical and least cost instruments are often used in the literature (Baum-Snow 2007; Duranton and Turner 2012; Faber 2014, Alder 2019).

model’s Stone-Geary preferences, since the rich are more likely to move into appreciating neighborhoods, given that they spend a smaller fraction of their income on housing.

4.4 Aggregate Effects from Reduced Form Sufficient Statistics

Table 6 measures TransMilenio’s aggregate effects by using the estimated reduced-form elasticities to implement the sufficient statistics approach outlined in Proposition 1.

First Order vs General Equilibrium Welfare Impacts. The standard approach to evaluate the gains from transit infrastructure is based on the “value of travel time savings” (e.g. Small and Verhoef 2007). Despite the rich channels captured in the general equilibrium model, Proposition 2 in Appendix C.5 shows that when the equilibrium is efficient, an application of the envelope theorem implies that this is precisely the first order welfare impact from a change in infrastructure.

Panel A of Table 6 simulates what Bogotá would have looked like in 2018 without TransMilenio, and then adds it back in under the different approaches.³⁶ The first column reports TransMilenio’s gains under the first order approximation or VTTS approach from Proposition 2. This delivers a welfare increase of 1.26%, accruing solely through time savings. The second column shows the welfare gains using the full model from Proposition 1 and the estimated elasticities. These deliver a much larger gain of 2.34%. The VTTS thus accounts for only 54% of the total welfare gains, yielding one of the paper’s central results—that equilibrium effects matter for valuing the gains from new transit infrastructure in cities. Confidence intervals for these main welfare effect are reported from a bootstrap procedure that accounts for the uncertainty in the model’s parameter estimates (see Appendix C.9 for details). While there is meaningful uncertainty surrounding these estimates, I can reject the null that the fraction of welfare gains accounted for by the VTTS is not less than one (p-value of 0.04). The difference between the equilibrium and the first order welfare effects could be due either to the size of the shock (since the approximation may perform poorly for large shocks) or deviation from efficiency (due to amenity and productivity externalities). The final column shows that when externalities are turned off, the VTTS explains a larger portion of the equilibrium effects. The size of the shock explains about one third of the 46% gap between the VTTS and the general equilibrium welfare effect, with the externalities accounting for the remaining two thirds.

Aggregate Effects. Panel B presents TransMilenio’s impact on aggregate outcomes using the results from Proposition 1.³⁷ Doing so requires values for α , β and σ in addition to the CMA elasticities. I estimate the $1 - \alpha = 0.206$ by computing the share of floorspace in total costs across establishments

³⁶I refer to the “2018 equilibrium” as the post-TransMilenio equilibrium. Population data and land market data come from 2018, employment data from 2015, land market data come from 2018, and the TransMilenio network includes all phases. Since there may be multiple equilibrium in the presence of externalities, the selection rule used is to start the algorithm from the observed equilibrium when solving for counterfactual equilibria. This can be rationalized through path dependence in a dynamic model of a city.

³⁷The percentage change in each variable under the counterfactual without the TransMilenio network are reported, i.e. $100 \times (X_{2015}^{NoTM} / X_{2015} - 1)$ for any variable X_{2015} . Numbers may therefore differ from Panel A which inverts the ordering by using the equilibrium without TransMilenio as the base. Table A.2 presents robustness of the welfare results in row 1 to alternative parameter values.

in each one digit non-agricultural industry, and averaging these by the sectoral employment shares in Bogotá.³⁸ I estimate $1 - \beta = 0.274$ from the average expenditure share of housing in Bogotá. Lastly I set $\sigma = 6$ close to median estimates from Feenstra et. al. (2018), but vary this in robustness checks.

The first column presents the closed city results from the model developed above. The second column presents results from an extension outlined in Appendix E.1, which allows for an upward-sloping supply of migrants into the city from the rest of the country. There are large aggregate impacts on welfare and city output under either mobility assumption. Without migration into Bogotá, GDP and welfare rise by 2.98% and 2.21% respectively, with a slight fall in the level of floorspace prices.³⁹ With migration, the welfare gain falls to 0.6% since the increase in population of 9.51% bid up floorspace values by 5.28%. GDP rises by 12.71%, but this is mostly due to population growth: GDP per capita rises by 5.67%. The final two rows show that TransMilenio can account for between 2.96% and 13.36% of Bogotá's GDP growth from 2000 to 2016, and up to 34.9% of observed population growth. TransMilenio's effects are quantitatively important, but not implausibly large. The third row shows that it was also a profitable investment for the city, leading to an increase of at least 2.5% in the steady state level of GDP net of construction and operating costs (see Appendix F.4 for details).

Incorporating Congestion. While speeds for cars and other types of buses did not change on routes adjacent to TransMilenio (see Appendix G), the BRT could have had aggregate effects on road speeds that do not appear in a difference-in-difference specification. Appendix E.2 extends the baseline model to gauge the impact of the BRT in the presence of congestion. The extension blends elements from Allen and Arkolakis (2021) and Gaduh et. al. (2022). The "economic module" of the model is unchanged: the same system of equations governs the response of economic activity to a change in commute times. However, a new "traffic module" is added that allows the change in commute times to depend on both new physical infrastructure and any changes in commuting patterns via congestion. The result is one combined system of equations where the change in economic activity and commute times are jointly determined in response to new infrastructure.

The results are shown in Panel C of Table 6, using the congestion elasticity of 0.06 estimated for Bogotá by Akbar and Duranton (2017). The first two rows show TransMilenio's welfare effect in the model with and without congestion.⁴⁰ Allowing for congestion leads to a larger welfare gain: as some commuters substitute away from cars onto the BRT, roads become less congested and driving times fall. This effect is small, however, with a welfare increase that is only 0.55% larger than without congestion. The last row assesses the welfare impact had the TransMilenio lanes been used to add new car instead of BRT lanes. The welfare effects are tiny in comparison: the welfare change would have been only 0.6% of the gains caused by TransMilenio. Overall, these results suggest that the baseline welfare effects provide a lower bound of the BRT system's impact in the presence of congestion.

³⁸The data on cost shares comes from the Encuesta Anual Manufacturera, Encuesta Anual de Servicios and the Encuesta Anual de Comercio in 2010. The sectoral employment shares are averages from 2000-2015 from the GEIH and ECH.

³⁹It is typical to have little change in the overall level of house prices in closed city models with fixed housing supply since the supply of housing and overall population is fixed.

⁴⁰The welfare effect in this model with the congestion elasticity set to zero differs from the baseline model, due to the different formulation of constructing commute costs.

5 Distributional Effects

While the sufficient statistics approach from the previous section can speak to the BRT’s aggregate effects, it is silent on the distribution of these impacts across worker groups. This section therefore estimates the full model from Section 3 to answer this question.⁴¹

5.1 Parameter Estimation

5.1.1 Parameters Estimated without Solving the Model

Externally Calibrated Parameters $\{\sigma, \sigma_D\}$ I set the elasticity of substitution between labor skill groups to $\sigma = 1/0.7$ based on the review in Card (2009), and $\sigma_D = 6$ as described in Section 4.4.⁴²

Share Parameters $\{\alpha_s, \beta, \alpha_{sg}\}$ I estimate $1 - \alpha_s = 0.206$ as described in Section 4.4 using data on the share of floorspace in total production costs, and set this to be equal across industries. I estimate $1 - \beta = 0.24$ to match Bogotá’s long-run housing expenditure share.⁴³ As described in Appendix D.2, the labor shares α_{sg} are calibrated to match the relative wage bill for college-educated workers in each industry.

Commute Costs and Elasticity The estimates κ, λ, b_m were provided in Table 1, delivering commute times t_{ija} under each car ownership status from (2) and (3). The commute elasticities θ_g can be estimated by taking logs and first differences of the expression for commute flows (4) to yield

$$\Delta \ln \pi_{j|ia} = \gamma_{ia} + \delta_{jg} - \theta_g \kappa \Delta t_{ija} + \varepsilon_{ija},$$

where γ_{ia} and δ_{jg} are fixed effects and ε_{ija} is an unobserved component of commute costs. Given a value of κ , θ_g is identified off the sensitivity of changes in commute flows to changes in commute times induced by TransMilenio for each group.⁴⁴

Table 7 presents the results. Across all specifications, high-skilled workers are found to have a lower θ_g (i.e. a higher dispersion of productivity shocks across locations) making their commute choices less sensitive to commute times. The overall magnitude and fact that more educated workers are estimated to have a greater dispersion of match-productivities line up with existing estimates (e.g. Lee 2020; Hsieh et. al. 2019; Galle et. al. 2022). Since the moment conditions in the following section use the instruments for TransMilenio’s placement, I use the IV estimates in column 2 as preferred estimates, which yield values of $\theta_H = 2.655$ and $\theta_L = 3.98$. I explore the robustness of the results to using the OLS or PPML estimates in Table A.4.

⁴¹A list of all parameters and sources of variation used to identify them is provided in Appendix A.3.

⁴²Estimating σ within the model would require a shock to relative labor supply and wage data by location of employment across space to measure the response in relative wages. Since the latter are not available, I calibrate σ instead and conduct robustness to both σ and σ_D in Table A.4.

⁴³See the Engel curves presented in Appendix G.

⁴⁴Another option here would have been to allow κ to vary by group. While this wouldn’t matter for this particular moment of the sensitivity of commute flows to commute times, allowing for this possibility in estimating the mode choice model led to $\kappa_H = 0.0126$ (0.0062) and $\kappa_L = 0.0113$ (0.0054). Since these are statistically indistinguishable from each other (p-value for a test of equality of 0.22), I use the assumption of common κ across groups.

5.1.2 Parameters Estimated Solving the Full Model

It remains to estimate the parameters $\{\bar{h}, p_a, T_g, \eta_g, \mu_A, \mu_{U,g}\}$. Appendix D.3 shows how, given prior parameter estimates, there is a vector $\{\bar{h}, p_a, T_g\}$ that matches the average expenditure share on housing, the average expenditure on cars, and the college wage premium, respectively.

The residential supply elasticity η_g and the spillover parameters $\mu_A, \mu_{U,g}$ are estimated via GMM. The intuition for identification is very similar to that of the sufficient statistics approach of Section 4. TransMilenio provides a shock to the attractiveness of each residential neighborhood through increased RCMA. The response of residential inflows to this shock identifies the residential supply elasticity. The response of model-inferred amenities to the resultant change in neighborhood composition identifies the amenity spillover. TransMilenio also provides a shock to the supply of workers commuting to each employment location through increased FCMA. The response of model-inferred productivity to this change in employment identifies the productivity spillover.

Amenities Moments Taking logs of the expression for resident supply (6) delivers

$$\Delta \ln L_{Riag} = \eta_g \Delta \ln V_{iag} + \eta_g \mu_{U,g} \Delta \ln \frac{L_{RiH}}{L_{Ri}} + \gamma_{\ell,g} + \gamma'_{R,g} \text{Controls}_i + \Delta \ln \epsilon_{Riag} \quad (20)$$

where $\Delta \ln V_{iag} \equiv \Delta \ln \tilde{y}_{iag} - (1 - \beta) \Delta \ln r_{Ri}$, $\gamma_{\ell,g}$ are locality-group fixed effects, and Controls_i denote tract characteristics (that can have separate effects by group) used to partially control for changes in unobserved amenities. $\Delta \ln \epsilon_{Riag}$ reflects residual variation in unobserved amenity growth.

The residential supply elasticity η_g is identified off the responsiveness of residential populations to exogenous variation in the common utility from living in a location $\Delta \ln V_{iag}$. This comes from my instruments for RCMA, which I use to construct predicted change in net income using the instrument for TransMilenio to generate expected income in the post-period.⁴⁵ Let $\Delta \ln \tilde{\Phi}_{Riag}^{IV}$ denote the expected growth in net income growth averaged across the LCP and Tram instruments (as in Table 5).

Identification of the spillovers $\mu_{U,g}$ requires exogenous variation in a neighborhood's college share. I use two instruments to this end. First, tracts that experience a greater growth in CMA to high-skilled jobs relative to low-skilled jobs should experience a larger increase in the share of college residents. This is captured by $Z_{Diff,i} = \Delta \ln \tilde{\Phi}_{RiH}^{IV} - \Delta \ln \tilde{\Phi}_{RiL}^{IV}$ where $\bar{X}_i = \sum_a X_{ia}$. Second, tracts with expensive housing where CMA improves should see a greater rise in the college share. This comes directly from log-linearizing the expression for residential populations (6). Intuitively, poor low-skilled residents are less willing to pay for increased access to jobs in expensive neighborhoods due to their greater expenditure on housing.⁴⁶ I capture this by interacting the change for

⁴⁵Letting $t - 1$ and t reference the pre- and post-TM periods, adjusted RCMA is $\tilde{\Phi}_{Riag,t-1} \equiv T_{g,t-1} \Phi_{Riag,t-1}^{1/\theta_g} - p_{a,t-1} a + \pi_{t-1}$ and $\tilde{\Phi}_{Riag,t}^{IV} \equiv T_{g,t} (\Phi_{Riag,t}^{IV})^{1/\theta_g} - p_{a,t} a + \pi_t$. The change is simply $\Delta \ln \tilde{\Phi}_{Riag}^{IV,k} = \ln \tilde{\Phi}_{Riag,t}^{IV,k} - \ln \tilde{\Phi}_{Riag,t-1}^{IV,k}$ for $k \in \{LCP, Tram\}$. Then $\Delta \ln \tilde{\Phi}_{Riag}^{IV} = E_k [\Delta \ln \tilde{\Phi}_{Riag}^{IV,k}]$.

⁴⁶Log-linearizing the expression for residential populations (6) yields

$$\Delta \ln L_{Riag} \approx \mu_{iag}^L \frac{\eta_g}{\theta_g} \Delta \ln \Phi_{Riag} - \eta_g (1 - \beta + \mu_{iag}^R) \Delta \ln r_{Ri} + \epsilon_{iag}$$

where $\epsilon_{iag} \equiv a \mu_{iag}^a \Delta \ln p_a + \mu_{iag}^\pi \Delta \ln \pi + \eta_g \Delta \ln u_{iag}$. Here $\mu_{iag}^L \equiv T_g \Phi_{Ri}^{1/\theta_g} / \tilde{y}_{iag}$ and $\mu_{iag}^R = r_{Ri} \bar{h} / \tilde{y}_{iag}$ are the share of labor

high-skilled residents with the house price in the initial period $Z_{Rents,i} = \Delta \ln \bar{\Phi}_{RiH}^{IV} \times \ln r_{Ri}^{2000}$.⁴⁷

The moment conditions used to identify η_g and $\mu_{U,g}$ are therefore⁴⁸

$$E[\Delta \ln \epsilon_{Riag} Z_{Riag}] = 0, \quad Z_{Riag} \in \left\{ \Delta \ln \bar{\Phi}_{Riag}^{IV}, Z_{Diff,i}, Z_{Rents,i} \right\}.$$

Productivity Moments Composite productivity $A_{js} \propto W_{js}^{\alpha_s} r_{Fj}^{1-\alpha_s} X_{js}^{1/(\sigma_D-1)}$ is the residual that ensures the model definition for sales holds. As shown in Proposition 3 in Appendix D.1, this can be recovered (to scale) using data on employment, residence, floorspace prices and commute costs. The model infers high productivity in locations where employment is high (reflected through high sales) relative to the observed price of commercial floorspace and the accessibility to workers through the commuting network (which determines wages). Taking logs of (14) and including a set of control variables to (partially) capture changing fundamentals yields

$$\Delta \ln A_{js} = \mu_A \Delta \ln \tilde{L}_{Fj} + \gamma_\ell + \gamma'_F \text{Controls}_j + \Delta \ln \epsilon_{Fjs}$$

where $\Delta \ln \epsilon_{Fjs}$ reflects residual variation in unobserved productivity growth.

The agglomeration elasticity is identified from the extent to which model-implied composite productivity depends on employment. Since employment will be correlated with unobserved components that make locations more productive, I use the instruments for FCMA growth as a labor supply shock. The moment conditions used to identify μ_A are therefore

$$E[\Delta \ln \epsilon_{Fis} Z_{Fig}] = 0, \quad Z_{Fig} \in \left\{ \Delta \ln \bar{\Phi}_{FiL}^{IV}, \Delta \ln \bar{\Phi}_{FiH}^{IV} \right\}.$$

Both sets of moments are stacked into a system of moment conditions which is estimated jointly in a single GMM estimation. I estimate standard errors via a block-bootstrap procedure, resampling at the tract-level to allow for arbitrary within-tract correlation in unobservables.⁴⁹

GMM Results Table 8 presents the main results. The productivity externality of 0.242 is slightly larger than existing estimates, although it is slightly noisy (p-value of 0.058), and thus contains smaller values within its confidence intervals. This is also one of the first estimates outside of a developed country setting. The residential population elasticity is slightly larger for the high-skilled than the low-skilled. The spillovers for residential amenities are 0.730 and 1.002 for low- and high-skilled workers. While both groups value living near high-types, the college educated value it most.

income and fixed housing expenditure of total net income. Note that μ_{iag}^R is greater for poor individuals since they spend a greater fraction of income on housing. Thus, poor low-skilled workers are more sensitive to house price appreciation and are less willing to pay for improved CMA than the high-skilled.

⁴⁷Controls for initial house prices are included to allow this characteristic to have its own impact on population growth; controls from the reduced form results are included and reported in Table 8.

⁴⁸Orthogonality conditions with each control variable are also included. The time periods used for pre- and post-periods for each variable are the same as in the previous section, and use the full 2013 TransMilenio network.

⁴⁹Bootstrapping is needed since units of observation vary across moment conditions, rendering the standard asymptotic variance formulas inapplicable. See Appendix D.4 for a benchmark of the amenity spillover estimates to Diamond (2016).

Model Validation The model’s fit of two non-targeted moments provides additional confidence in its results. First, Figure A.4 plots the observed change in the share of floorspace used for residential purposes against that predicted by the model. While the two correlate well, the correct test of the model is not that the correlation or R2 is high but rather that the regression of the observed on the model-predicted changes has a slope of one.⁵⁰ The slope of this regression is 1.598 (0.822), which is statistically indistinguishable from one (p-value 0.47). Second, the model predicts that changes in income are related to RCMA through $d \ln \bar{y}_i = \frac{1}{\theta} d \ln \Phi_{Ri}$ with elasticity $1/\theta$. This regression is reported in column (3) of Table A.16. The coefficient of 0.522 (0.224) is statistically indistinguishable from $1/\theta$ after plugging in the estimate of $\theta = 3.398$ (p-value 0.31).

5.2 Results

Panel A of Table 9 presents the main result: welfare inequality increases by 0.55% as a result of TransMilenio. It should be noted that the confidence intervals convey uncertainty in this estimate, and the test of whether the high-skill gain more than the low-skill only has a p-value of 0.15. With this caveat in mind, I turn understanding the source of this result.

Why would the high-skilled benefit the most? Panel B decomposes the welfare gains, starting with a simplified case of the model and slowly adding its ingredients to isolate each one’s impact.

The first row assumes workers share the same (average) value for η and θ and are perfect substitutes in production. This model is similar to the simple model used in Section 4.4 since it abstracts from heterogeneity across workers. Reassuringly, the average welfare effect of 2.19% is very close to the 2.335% reported from the sufficient statistics approach in Table 6. Low-skilled workers benefit the most, with inequality falling by 0.37%.

The second row allows workers to differ in their commuting elasticities. Recall that the high-skilled have a lower commute elasticity. This shifts the gains towards the high-skilled, with inequality now falling by 0.16%. A lower commuting elasticity tends to increase the incidence of high commute costs, since workers have very sticky preferences for workplace locations and are less able to substitute away to less costly options when transit is slow. The third row allows the residential choice parameters to equal their estimated values, with a modest reduction in the fall in inequality.

The last thing that changes as one moves to the result from the full model in Panel A is that workers are imperfect substitutes in production. Intuitively, the average welfare effect falls. For example, a large inflow of low-skilled workers will increase the supply of the labor bundle less than when both types are perfect substitutes. However, this also causes the sign of the impact on inequality to switch, with welfare inequality rising by 0.55%. This occurs for two reasons. First, high-skilled workers are now partially shielded from the reduction in wages due to the large labor supply shift of low-skilled workers who use public transit since they are no longer perfect substitutes. Second, it now matters whether each skill group is connected to locations where demand for their skill is

⁵⁰Other shocks orthogonal to the model may cause the correlation or R2 between observed and model-predicted changes to fall. The model is trying to capture the counterfactual of how activity would have changed if the only shock had been the change in infrastructure. For discussion see Adao, Costinot and Donaldson (2023).

highest. For the geography of Bogotá and TransMilenio, this tends to benefit the high-skilled who are concentrated in the city's north which TransMilenio connected with the high skill-intensive industries in the center and center-north. Residence and employment for the low-skilled is more dispersed, so TransMilenio connected a smaller fraction of these workers with high-wage locations.

Overall, these results suggest that the incidence of improving public transit depends not only on how much each group uses it, but also on how willing each group is to bear high commute costs to work at a particular location, whether the system connects worker groups with their high-wage locations and the general equilibrium response of wages and house prices. In the context of Bogotá's TransMilenio, these reallocation and equilibrium effects are large enough to reverse the effects on inequality, which ultimately rose 0.55% as a result of the BRT.

Domestic Services and Alternative Home Ownership Assumptions. From 2000-2014, 7.3% of non-college educated Bogotanos worked as domestic helpers while almost no college-educated workers did. On the one hand, the model may underestimate the gains to the low-skilled by ignoring the fact that TransMilenio improved access to domestic services jobs in the homes of the college educated in the North. On the other hand, the high-skilled also benefitted from this increased labor supply, which lowered the cost of hiring domestic workers. Appendix E.5 extends the model to incorporate employment in domestic services, and Panel C of Table 9 presents the results. Overall, these two effects tend to balance out—the increase in inequality is very similar to the main model in Panel A. The last two rows of Panel C incorporate different assumptions over home ownership as outlined in Appendix E.6, with the results fairly invariant across the alternatives.⁵¹

5.3 Policy Counterfactuals

Impact of Alternative Networks. The first panel of Table 10 reports the impact on welfare, inequality, and output had the network been built without lines A and H, which connect the city's north and south with the CBD. The line to the south has a bigger effect on welfare (which would have been 0.3% lower without it), which is logical given the higher population density of poor and middle income workers there. For the same reason, the line to the north has the greater effect on inequality (which would have risen by 0.2% less without it). Intuitively, each group benefits relatively more from lines that improve accessibility from where they live.

The welfare gains from these trunk lines are exceeded by the benefits from the feeder bus network (as welfare would have been 0.94% lower without it). These buses connect outlying areas with portals and run on existing roadways. By providing complementary services that reach residents in outlying but dense residential areas, they can solve the last-mile problem of traveling between stations and final destinations. Given the low cost of feeder systems compared to the capital-intense BRT, these results suggest a high return to policy makers considering cheap, complementary services to increase access to mass rapid transit infrastructure.

⁵¹Table A.4 reports robustness of the main results to (i) allowing migration into the city, (ii) a larger elasticity of substitution between labor types, (iii) alternative employment data, (iv) alternative elasticities of demand, (v) alternative commuting elasticities and (iv) a decision over where to live and work.

A key trade-off policymakers face is whether to build fast rail, medium-speed BRT, or slower bus networks. The speed of such networks could affect the distributional consequences, for example, if the high-skilled were especially willing to pay to live near faster networks and priced out poorer residents. The last row runs a counterfactual that increases TransMilenio's speed to 35 km/h, close to the average operating speed of London's Underground. The increase in both welfare and inequality would have been much higher, confirming the intuition that faster systems benefit the rich relatively more. However, Figure A.5 compares the change in the college share near stations under this counterfactual. While it does increase in tracts closer than 500m from a station, the increase is very mild. This suggests the channels mentioned above, rather than gentrification, are responsible for why the rich benefit more from faster transit.

Land Value Capture One main criticism of TransMilenio was that its construction was not accompanied by an adjustment of zoning laws to allow housing supply to respond where it was needed. Appendix G shows that housing supply did not respond to the system's construction, consistent with other evidence on the restrictive role played by land use regulation (Cervero et. al. 2013). Many cities, such as Hong Kong and Tokyo, have had success in implementing LVC schemes which increase permitted densities around new stations but charge developers for the right to build there (Hong et. al. 2015). These policies increase housing supply and raise revenue to finance the infrastructure's construction.

I now evaluate the impact of TransMilenio had housing supply responded to the opening of the system. As a benchmark, I allow housing supply to adjust to the increase in floorspace values following a log-linear supply curve. Given that I do not observe a housing supply response in Bogotá that would permit me to measure a city-specific housing elasticity, I instead make a conservative choice and assume the housing supply in Bogotá is the same as that in Oakland, CA, the 6th most inelastic city in the US according to Saiz (2010). I then simulate the effect of two potential LVC schemes. First, I assume the government sells the rights to developers to increase floorspace by a maximum of 30% in tracts within 500m of stations, mimicking the "development rights sales" undertaken in certain Asian, European, and American cities.⁵² Second, I assume the government sells permits that allow for the same change in total floorspace, but instead allocates the permitted floorspace changes according to a location's predicted change in CMA. Details on this model extension are provided in Appendix E.4. I compare the two equilibria by first removing TransMilenio (without housing adjustment) and then by adding it back under each housing supply model.

The last two panels of Table 10 present the results. Panel B shows the impacts on welfare. Under free adjustment, welfare would have been 44.04% higher than it is today. Under the LVC schemes, welfare would have been 24.47% or 43.82% higher than it is today under the distance- or CMA-based policies respectively (with similar relative effects on city output). These welfare improvements come from increasing housing supply where it is demanded the most as a result of new infrastructure,

⁵²See Hong et. al. (2015) and Salon (2014) for further details. The parameters of this counterfactual are motivated by the example of Nanchang, China, where floor area ratios were increased by a uniform amount within 500m of stations. Revenues from the scheme covered 20.5% of costs, similar to my results.

tempering down floorspace price appreciation. The high return to the CMA-based instrument highlights how well-targeted zoning adjustments that allocate permits towards where they are needed most deliver bigger benefits. Panel C shows the fiscal benefits of the LVC schemes. Depending on how much the city population grows in response to the BRT, the distance-based instrument recoups 4-11% of construction costs, while the CMA-based scheme covers 6-21% of such costs.

These results suggest the potential for large welfare gains to governments pursuing a unified transit and land use policy. These policies can also be used to finance the construction of public transit, and targeting zoning adjustments based on where demand for housing will increase the most delivers the largest benefits.

6 Conclusion

This paper makes three contributions to our understanding of the aggregate and distributional effects of urban transit systems. First, it develops a sufficient statistics approach to evaluate the aggregate effects of new transit infrastructure in cities. Second, it shows that these statistics can be measured from readily available data and estimates them using the variation in accessibility induced by TransMilenio's construction. Third, it quantifies the welfare gains from the BRT under the equilibrium model and compares these gains with the VTTS to isolate the importance of reallocation and general equilibrium effects. Fourth, it estimates a richer model to that nests the sufficient statistics approach to quantify who the gains are shared between the rich and poor.

The study finds that the quantitative urban model performs well in explaining the adjustment of economic activity to transit infrastructure, with the log-linear relationships predicted by the model borne out in the data. The VTTS only account for around 57.5% of the total welfare gain from the new transit infrastructure. Thus, accounting for equilibrium effects matters for valuing the gains from new transit infrastructure in cities, and the framework developed in this paper provides a blueprint to do so. It also finds that the accounting for reallocation and general equilibrium effects acts against the benefits to poor workers who tend to use transit the most, which in the case of TransMilenio meant that welfare inequality rose by a mild 0.55%.

The paper also provides two key insights that can inform transit infrastructure policy. The first is that low-cost "feeder" bus systems that complement mass rapid transit by providing "last-mile" service for passengers' easy access to a system's terminals yield high returns. The second is that the welfare gains would have been around 40% higher had the the government implemented a more accommodative zoning policy, and government revenues from an LVC scheme could have raised a significant portion of construction costs. This underscores the benefits to cities from pursuing a unified transit and land use policy.

References

- ADAO, R., ARKOLAKIS, C., AND F. ESPOSITO (2019), "Spatial Linkages, Global Shocks, and Labor Markets: Theory and Evidence," Working Paper.
- ADAO, R., COSTINOT, A., AND D. DONALDSON (2023), "Putting Quantitative Models to the Test: An Application to Trump's Trade War," Working Paper.
- AHLFELDT, G.M., REDDING, S.J., STURM, D.M. AND N. WOLF (2015), "The Economics of Density: Evidence from the Berlin Wall," *Econometrica*, 83(6): 2127-2189.
- AKBAR, P., AND G. DURANTON (2017), "Measuring the Cost of Congestion in a Highly Congested City: Bogotá," Working Paper.
- AKBAR, P., G. DURANTON, COUTURE, V., GHANI, E., AND A. STOREYGARD (2021), "Accessibility and Mobility in Urban India," Working Paper.
- ALCALDÍA MAYOR DE BOGOTÁ D.C. (2009), "Evaluación de la Alternativas de la Red Metro Del SITP," <https://www.metrodebogota.gov.co/sites/default/files/documentos/Producto%2015.%20Tomo%201.%20F>
- ALDER, S., (2019), "Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development," Working Paper.
- ALLEN, T., ARKOLAKIS, C. AND X. LI (2014), "On the Existence and Uniqueness of Trade Equilibria," Working Paper.
- ALLEN, T., ARKOLAKIS, C. AND X. LI (2015), "Optimal City Structure," Working Paper.
- ALLEN, T., AND C. ARKOLAKIS (2019), "The Welfare Effects of Transportation Infrastructure Improvements," NBER Working Paper 25487.
- ALLEN, T., AND C. ARKOLAKIS (2021), "The Welfare Effects of Transportation Infrastructure Improvements," Working Paper.
- ANSELIN, L., AND N. LOZANO-GRACIA (2012), "Is the price right?: Assessing estimates of cadastral values for Bogotá, Colombia," *Regional Science Policy & Practice*, 4(4): 495-508.
- BALTES, M., J. BARRIOS, A. CAIN, G. DARIDO, P. RODRIGUEZ (2006), "Applicability of Bogotá's Trans-Milenio BRT System to the United States," Federal Transit Administration Report Number FL-26-7104-01.
- BARR, J., J. BEVERIDGE, C. CLAYTON, A. DANAHAR, J. GONSALVES, B. KOZIOL, S. RATHWELL (2010), "Designing Bus Rapid Transit Running Ways", American Public Transportation Association, Recommended Practice Report APTA-BTS-BRT-RP-003-10.
- BARTELME, D. (2018), "Trade Costs and Economic Geography: Evidence from the U.S.," Working Paper.
- BAUM-SNOW, N. (2007), "Did Highways Cause Suburbanization?" *Quarterly Journal of Economics*, 122: 775-805.
- BAUM-SNOW, N. , AND M. KAHN (2005), "Effects of Urban Rail Transit Expansions: Evidence from Sixteen Cities, 1970-2000," *Brookings-Wharton Papers on Urban Affairs*, 147-206.
- BAUM-SNOW, N., BRANDT, L., HENDERSON, J., TURNER, M., AND Q. ZHANG (2017), "Roads, Railroads and Decentralization of Chinese Cities," *Review of Economics and Statistics*, 99(3): 435-448.
- BAYER, P., FERREIRA, F., AND R. MCMILLAN (2007), "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 115(4): 588— 638.
- BILLINGS, S. (2011), "Estimating the value of a new transit option," *Regional Science and Urban Economics*, 41(6):525-536.
- BIRD, J AND A. VENABLES. (2019), "The efficiency of land-use in a developing city: traditional vs modern tenure systems in Kampala, Uganda," *CEPR Discussion Papers 13563, C.E.P.R. Discussion Papers*.
- BOCAREJO, J-P., PEREZ, M.A., AND I. PORTILLA (2013), "Impact of Transmilenio on Density, Land Use,

- and Land Value in Bogotá," *Research in Transportation Economics*, 40:78-86.
- BORUSYAK, K., AND P. HULL (2021), "Non-Random Exposure to Exogeneous Shocks," Working Paper.
- BRT DATA (2017), *Global BRT Statistics*, Retrieved from <http://brtdata.org>.
- BRYAN, G., AND M. MORTEN (2019), "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127(5):2229-2268.
- CARD, D. (2009), "Immigration and Inequality," *American Economic Review*, 99(2): 1–21.
- CERVERO, R., IUCHI, K., AND H. SUZUKI (2013), *Transforming Cities with Transit: Transit and Land-Use Integration for Sustainable Urban Development.*, Washington, DC: World Bank.
- COMBES, P.-P., G. DURANTON, L. GOBILLON, AND S. ROUX (2010), "Estimating Agglomeration Economies with History, Geology, and Worker Effects," in *Agglomeration Economies*, ed. by E. L. Glaeser, Chicago University Press, Chicago.
- CONLEY, T. (1999), "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 92(1), 1–45.
- CRACKNELL, J. (2003), "Transmilenio Busway-Based Mass Transit, Bogotá, Colombia," World Bank Working Paper.
- DAVIS, M., AND F. ORTALO-MAGNÉ (2011), "Household expenditures, wages, rents," *Review of Economic Dynamics*, 14(2):248-261.
- DIAMOND, R. (2016), "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 106(3):479–524.
- DONALDSON, D. (2015), "The Gains from Market Integration," *Annual Review of Economics*, 7:619-647.
- DONALDSON, D., AND R. HORNBECK (2016), "Railroads and American Economic Growth: A Market Access Approach," *Quarterly Journal of Economics*, 131(2): 799-858.
- DONALDSON, D. (2018), "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, 108(4-5):899-934.
- EATON, J., AND S. KORTUM (2002), "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.
- FAJGELBAUM, P. D., AND E. SCHAAL (2020), "Optimal Transport Networks in Spatial Equilibrium," *Econometrica*, 88(4):1411-1452.
- FEENSTRA, R., P. LUCK, M. OBSTFELD AND K. RUSS (2018), "In Search of the Armington Elasticity," *The Review of Economics and Statistics*, 100(1): 135-150.
- FERNANDEZ-VAL, I. AND M. WEIDNER (2016), "Individual and Time Effects in Nonlinear Panel Data Models with Large N, T," *Journal of Econometrics*, 192(1): 291-312
- FUJIMOTO, T., AND U. KRAUSE (1985), "Strong ergodicity for strictly increasing nonlinear operators," *LINEAR ALGEBRA AND ITS APPLICATIONS*, 71: 101-112.
- GADUH, A., GRACNER, T., AND A ROTHENBERG (2022), "Life in the Slow Lane: Unintended Consequences of Public Transit in Jakarta," *Journal of Urban Economics*, 128.
- GALLE, S., RODRIGUEZ-CLARE, A., AND M. YI (2022), "Slicing the Pie: Quantifying the Aggregate and Distributional Consequences of Trade," *The Review of Economic Studies*.
- GANONG P., AND D. SHOAG (2017), "Why has regional income convergence in the U.S. declined?" *Journal of Urban Economics*, 102:76-90.
- GIANNONE, E. (2021), "Skill-Biased Technical Change and Regional Convergence," Working Paper.
- GIBBONS, S., AND S. MACHIN (2005), "Valuing rail access using transport innovations," *Journal of Urban Economics*, 57(1):148–169.
- GLAESER, E., KAHN, M., AND J. RAPPAPORT (2008), "Why do the poor live in cities? The role of public transportation," *Journal of Urban Economics*, 63(1):1–24.

- GONZALEZ-NAVARRO, M., AND M. TURNER (2018), "Subways and urban growth: Evidence from Earth," *Journal of Urban Economics*, 108:85-106.
- GUERRIERI, G., HARTLEY, D., AND E. HURST (2013), "Endogenous Gentrification and Housing Price Dynamics," *Journal of Public Economics*, 100:45-60.
- GREENSTONE, M., HORNBECK, R., AND E. MORETTI (2010), "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings," *Journal of Political Economy*, 118(3), 536–598.
- HEBLICH, S., REDDING, S., AND D. STURM (2020), "The Making of the Modern Metropolis: Evidence from London," *The Quarterly Journal of Economics*, 135(4):2059-2133.
- HIDALGO, D., AND P. GRAFTIEAUX (2008), "Bus Rapid Transit Systems in Latin America and Asia: Results and Difficulties in 11 Cities," *Journal of the Transportation Research Board*, 2072(1):77-88.
- "HISTORIA DE TRANSMILENIO" TRANSMILENO.GOV.CO. FEB 15, 2022. [HTTPS://WWW.TRANSMILENIO.GOV.CO/PUNTO-DE-TRANSMILENIO/](https://www.transmilenio.gov.co/PUNTO-DE-TRANSMILENIO/).
- HONG, Y-H., MURAKAMI, J., SUZUKI, H., AND B. TAMAYOSE (2015), *Financing Transit-Oriented Development with Land Values*, Washinton, DC: The World Bank.
- HSIEH, C-T., HURST, E., JONES, C., AND P. KLENOW (2019), "The Allocation of Talent and US Economic Growth," *Econometrica*, 87(5):1439-1474.
- KLINE, P., E. MORETTI (2014), "Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *Quarterly Journal of Economics*, 129(1), 275–331.
- LAWELL, C., UMANSKAYA, V., AND ZHANG, W., (2017), "The Effects of License Plate-Based Driving Restrictions on Air Quality: Theory and Empirical Evidence," *Journal of Environmental Economics and Management*, 82:181-220.
- LEE, E. (2020), "Trade, Inequality, and the Endogenous Sorting of Heterogeneous Workers," *Journal of International Economics*, 125.
- MACKIE, P., J. NELLTHORP, AND J. LAIRD (2005), "Notes on the Economic Evaluation of Transport Projects," World Bank Transport Note No. TRN-5.
- MCFADDEN, D. (1974), "The measurement of urban travel demand", *Journal of Public Economics*, 3(4): 303-328.
- MCKINSEY (2016), *Bridging Global Infrastructure Gaps*, McKinsey Global Institute
- MENCKHOFF, G. (2005), "Latin American Experience with Bus Rapid Transit," presented at Institute of Transportation Engineers Annual Meeting, Melbourne, Australia.
- MONTE, F., REDDING, S., AND E. ROSSI-HANSBERG (2018), "Commuting, Migration and Local Employment Elasticities," *American Economic Review*, 108(12):3855-3890.
- MORRISON, A., (2007), "The Tramways of Bogotá, Colombia," <http://www.tramz.com/co/bg/ft/te.html>
- MÜLLER, J.M. (2014), Bus Rapid Transit: The Answer to Transport Problems in Megacities? The Example of TransMilenio (Bogotá, Colombia), in Kraas, F., Aggarwal, S., Coy, M., Mertins, G. (eds) Megacities, International Year of Planet Earth., Springer, Dordrecht.
- OWENS, R., E. ROSSI-HANSBERG, AND P-D. SARTE (2020), "Rethinking Detroit," *American Economic Journal: Economic Policy*, 12(2):258-305.
- REDDING, S., AND D. STURM (2008), "The Costs of Remoteness: Evidence from German Division and Reunification," *American Economic Review*, 98(5), 1766–1797.
- ROSENTHAL, S., AND W. STRANGE (2004), "Evidence on the Nature and Sources of Agglomeration Economics," in *Handbook of Regional and Urban Economics*, ed. by J. V. Henderson, and J. Thisse, vol. 4. Elsevier North Holland, Amsterdam.

- ROSSI-HANSBERG, E., P. SARTE, AND R. OWENS III (2010), "Housing Externalities," *Journal of Political Economy*, 118(3): 485-535.
- RUIZ, F., AND G. VALLEJO (2010), "Using Land Registration as a Tool to Generate Municipal Revenue: Lessons from Bogota," *Cadastral Working Paper*.
- SAIZ, A. (2010), "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125(3):1253-1296.
- SALON, D. (2014), "Location Value Capture Opportunities for Urban Public Transport Finance," *Transit Leadership Summit, London, White Paper*.
- SEVEREN, C. (2021), "Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification," *The Review of Economics and Statistics*.
- STEINER, R., AND H. VALLEJO (2010), *Economy, in Colombia: a country study*, Library of Congress. Federal Research Division.
- SMALL, K., AND E. VERHOEF (2007), *The Economics of Urban Transportation*, New York: Routledge.
- TRAIN, K., AND D. MCFADDEN (1978), "The goods/leisure tradeoff and disaggregate work trip mode choice models," *Transportation Research*, 12(5): 349-353.
- TRANSPORTATION REVIEW BOARD (2003), "Bogotá, Colombia. Brief: TransMilenio BRT", *Transit Cooperative Research Program, Report 90*.
- TRANSPORTATION REVIEW BOARD (2007), "Bus Rapid Transit Practitioner's Guide", *Transit Cooperative Research Program, Report 118*.
- URIBE SANCHEZ, M.C. (2010), "Land Information Updating, a De Facto Tax Reform: Bringing up to Date the Cadastral Database of Bogota," *Cadastral Working Paper*.
- WORLD BANK (2009), *2009 Annual Report: Energy Sector Management Assistance Program (ESMAP)*, Washington, DC: World Bank.
- WRIGHT, L., AND W. HOOK (2007), *Bus Rapid Transit Planning Guide*, New York: Institute for Transportation and Development Policy.

Tables

Table 1: Mode Choice and Commuting Parameter Estimates

Panel A: Mode Choice	
Parameter	Estimate
κ	0.012** (0.006)
b_{Bus}	-0.085* (0.051)
b_{Car}	0.853*** (0.291)
b_{TM}	-0.212* (0.108)
λ	0.138** (0.067)
N	19,510

Panel B: Commute Semi-Elasticity (Aggregate)	
Parameter	Estimate
θ_{κ}	0.039** (0.016)
N	710

Note: Panel A shows estimation results from nested logit regression on mode choices from trip-level data from the 2015 Mobility Survey. Controls for hour of trip departure dummies and dummies for gender and quintiles of the age distribution are included for each mode, which is equivalent to allowing preferences for each mode to vary by these characteristics. Heteroscedasticity robust standard errors are reported in parentheses. Panel B shows gravity equation estimation results, estimated via PPML. The outcome is the log number of commuters between each origin and destination locality pair in 1995 or 2015. Fixed effects for each origin-destination pair, origin-year and destination-year pair are included. Reported coefficient is that on travel time. In both panels, only trips to work during rush hour (5-8am) by heads of households included. Controls of origin-destination pair characteristics interacted with year dummies include (i) the average number of crimes per year from 2007-2014, (ii) the average log house price in 2012 and (iii) the share of the trip that takes place along a primary road along the least-cost routes between origin and destination. Standard errors are clustered at the origin-destination locality. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 2: Baseline Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Residents							
ln(Res Floorspace Price)	0.370** (0.179)	0.382** (0.172)	0.384** (0.169)	0.342* (0.174)	0.379*** (0.147)	0.228** (0.090)	0.365* (0.202)
<i>N</i>	2,201	2,201	2,201	2,166	2,199	2,201	2,201
<i>R</i> ²	0.41	0.43	0.43	0.43	0.43	0.43	0.43
ln(Res Population)	0.722** (0.337)	0.648* (0.337)	0.746** (0.331)	0.630* (0.334)	0.683** (0.293)	0.318* (0.174)	1.086*** (0.388)
<i>N</i>	2,256	2,256	2,256	2,219	2,255	2,254	2,256
<i>R</i> ²	0.34	0.35	0.37	0.37	0.37	0.37	0.37
Panel B: Firms							
ln(Comm Floorspace Price)	0.526** (0.253)	0.540** (0.251)	0.621** (0.248)	0.580** (0.249)	0.514** (0.214)	0.362*** (0.135)	0.718** (0.301)
<i>N</i>	2,080	2,080	2,080	2,047	2,089	2,083	2,080
<i>R</i> ²	0.09	0.10	0.11	0.11	0.11	0.11	0.11
Comm Floorspace Share	0.290*** (0.087)	0.297*** (0.088)	0.291*** (0.088)	0.284*** (0.088)	0.197*** (0.072)	0.151*** (0.047)	0.286*** (0.101)
<i>N</i>	2,230	2,230	2,230	2,195	2,239	2,233	2,230
<i>R</i> ²	0.14	0.15	0.15	0.15	0.15	0.15	0.15
ln(Establishments)	2.101*** (0.735)	1.787** (0.761)	1.329* (0.751)	1.283* (0.757)	1.414** (0.642)	0.924** (0.395)	1.414 (0.888)
<i>N</i>	2,028	2,028	2,028	1,995	2,028	2,028	2,028
<i>R</i> ²	0.65	0.67	0.68	0.68	0.68	0.68	0.68
Locality FE	X	X	X	X	X	X	X
Log Dist CBD X Region FE	X	X	X	X	X	X	X
Basic Tract Controls	X	X	X	X	X	X	X
Historical Controls		X	X	X	X	X	X
Land Market Controls			X	X	X	X	X
Exclude Portals+CBD				X			
Exclude Band					1.5km		
Alt Time Aggregation						X	
Distance to TM Controls							X

Note: Observation is a census tract. Each entry reports the coefficient from a regression of the change in the variable in each row on the change in firm or residential commuter market access (RCMA for residential outcomes, FCMA for commercial outcomes). CMA is always computed holding employment and population fixed at their initial levels and excluding the location itself from the summation. Each column corresponds to a specification. In land market regressions of row 1, 3 and 4, outcomes are log changes between 2018 and 2000 and the change in CMA is that induced by all three phases holding residential population and employment fixed at their levels in 2000 (population in 2000 is a linear interpolation from the 1993 and 2005 census; employment is from the 2000 CCB data). In population regressions of row 2, the outcome is the log change in residential population between 1993 and 2018. The change in CMA is that induced by all three phases of TransMilenio, holding residential population and employment fixed at their levels in 1993 and 1990 respectively (measured from the population and economic censuses). In establishment regressions of row 5, the outcome is the log change in the number of establishments between 2000 and 2015 from the CCB data against the same CMA measures as the land market regressions. Establishment specifications are weighted by the share of establishments in a tract in the initial period. CBD X Region controls are log distance to the CBD, interacted with dummies for whether the locality is in the North, West or South of the city. Basic tract controls include (i) log area, (ii) log distance to the main road, (iii) log distance to a main road interacted with log distance to the CBD, (iv) dummies for each quartile of 1993 population density, 1990 employment share (employment divided by employment plus population), and 1993 college share. Historical controls include dummies for each quartile of population density in 1918, and a dummy for whether the tract was closer than 500m to a main road in 1933. Land market controls include the share of land developed, floor area ratio, share of floorspace used for commercial purpose, and log average floorspace value in 2000. Any control that represents the initial value of an outcome variable is dropped from that specification. Columns (1) to (3) incrementally add controls. Column (4) restricts the sample to tracts more than 500m from a portal or the CBD. Column (5) computes the change in market access to tracts further than 1.5km from the tract itself. Column (6) assumes users take the quickest form of public transit (i.e. the minimum rather than the weighted average within the public nest). Column (7) includes a dummy for whether tract is closer than 500m from any TransMilenio station. Heteroscedasticity and Autocorrelation Consistent (HAC) standard errors (Conley (1999)) with a 0.5km bandwidth reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 3: Planned Networks

	(1)	(2)
Panel A: Residents		
ln(Res Floorspace Price)		
$\Delta \ln(\text{CMA})$	0.384** (0.169)	0.368** (0.175)
$E[\Delta \ln(\text{CMA Plan})]$		0.084 (0.225)
N	4,402	4,402
R^2	0.43	0.43
p-val		0.71
ln(Res Population)		
$\Delta \ln(\text{CMA})$	0.746** (0.331)	0.817** (0.337)
$E[\Delta \ln(\text{CMA Plan})]$		-0.388 (0.371)
N	4,512	4,512
R^2	0.37	0.37
p-val		0.29
Panel B: Firms		
ln(Comm Floorspace Price)		
$\Delta \ln(\text{CMA})$	0.621** (0.248)	0.687*** (0.259)
$E[\Delta \ln(\text{CMA Plan})]$		-0.360 (0.426)
N	4,160	4,160
R^2	0.11	0.11
p-val		0.39
Comm Floorspace Share		
$\Delta \ln(\text{CMA})$	0.291*** (0.088)	0.290*** (0.089)
$E[\Delta \ln(\text{CMA Plan})]$		0.005 (0.099)
N	4,460	4,460
R^2	0.15	0.15
p-val		0.96
ln(Establishments)		
$\Delta \ln(\text{CMA})$	1.329* (0.751)	1.170 (0.778)
$E[\Delta \ln(\text{CMA Plan})]$		0.829 (0.692)
N	4,056	4,056
R^2	0.68	0.68
p-val		0.24

Note: Column (1) repeats the baseline specification i.e. column (3) from Table 2. That is, each entry reports the coefficient from a regression of the change in the variable in each row on the change in firm or residential commuter market access (RCMA for residential outcomes, FCMA for commercial outcomes). Column (2) adds as an additional explanatory variable the average change in RCMA or FCMA (depending on the outcome, RCMA for residential and FCMA for commercial) each tract would have received had TransMilenio been built across the 4 historical plans. The p-value corresponds to a χ^2 test of equality of coefficients on $\Delta \ln(\text{CMA})$ in columns 1 and 2. Heteroscedasticity and Autocorrelation Consistent (HAC) standard errors (Conley (1999)) with a 0.5km bandwidth reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 4: Staggered Station Openings

	(1)	(2)	(3)	(4)
Panel A: Residents				
	ResPr	ResPr	ResPop	ResPop
ln(RCMA)	0.183* (0.101)	0.231** (0.109)	0.396* (0.237)	0.521** (0.253)
ln(RCMA) Later Phase		0.336 (0.239)		0.345 (0.348)
<i>N</i>	2,144	2,144	2,207	2,207
<i>R</i> ²	0.45	0.45	0.29	0.29
Panel B: Firms (Land Markets)				
	CommPr	CommPr	CommSh	CommSh
ln(FCMA)	0.483** (0.199)	0.478** (0.201)	0.210*** (0.057)	0.206*** (0.056)
ln(FCMA) Later Phase		0.281 (0.666)		0.220 (0.172)
<i>N</i>	2,055	2,055	2,182	2,182
<i>R</i> ²	0.06	0.06	0.08	0.08
Panel C: Firms (Census Employment)				
	Emp	Emp	Form Emp	Form Emp
ln(FCMA)	1.640 (1.064)	1.791* (1.071)	2.097 (1.339)	2.143 (1.348)
ln(FCMA) Later Phase		1.497 (1.288)		0.438 (1.745)
<i>N</i>	1,927	1,927	1,629	1,629
<i>R</i> ²	0.23	0.23	0.17	0.17

Note: Table repeats the baseline specification i.e. column (3) from Table 2. Outcomes are (growth in) residential floorspace prices (Res Pr), residential population (Res Pop), commercial floorspace prices (Comm Pr), commercial floorspace share (Comm Sh), employment from the census (Emp), employment in establishments with more than 10 workers (Form Emp). For land market outcomes, the change in outcomes are measured between 2008 and 2000. The right hand side variables include CMA growth due to (i) phases 1 and 2 of the system open by 2006 (ln(*CMA*)) and (ii) phase 3 of the system open by 2013 (ln(*CMA*) Later Phase). For residential population, the change in outcome is measured between the 2005 and 1993 census. The right hand side variables include CMA growth due to phase 1 (open by 2003, with most opening by 2001), and the change in CMA due to phases 2 and 3 (opened in 2006 and 2013). For employment, the change in employment is measured from between the 2005 and 1990 economic censuses. The CMA variables are the same as for residential population. Heteroscedasticity and Autocorrelation Consistent (HAC) standard errors (Conley (1999)) with a 0.5km bandwidth reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: IV Estimates

	Baseline	IV	IV-LCP&Tram Exc Own	IV-LCP&Tram Exc 1.5km
Panel A: Residents				
ln(Res Floorspace Price)	0.384** (0.169)	0.276*** (0.102)	1.117*** (0.247)	1.111*** (0.229)
<i>N</i>	2,201	2,201	2,202	2,202
F-Stat		2,475.45	166.46	201.21
ln(Res Population)	0.746** (0.331)	0.553** (0.225)	0.554 (0.483)	0.653 (0.438)
<i>N</i>	2,256	2,256	2,239	2,239
F-Stat		2,404.19	248.56	295.97
Panel B: Firms				
ln(Comm Floorspace Price)	0.621** (0.248)	0.552*** (0.204)	0.397 (0.291)	0.607** (0.304)
<i>N</i>	2,080	2,080	2,085	2,085
F-Stat		3,165.71	664.43	746.85
Comm Floorspace Share	0.283*** (0.093)	0.257*** (0.071)	0.227** (0.105)	0.204** (0.103)
<i>N</i>	2,231	2,230	2,235	2,235
F-Stat		3,112.12	670.54	733.74
ln(Establishments)	1.329* (0.751)	1.229** (0.562)	2.207** (0.880)	1.954** (0.827)
<i>N</i>	2,028	2,028	1,995	1,995
F-Stat		2,878.76	402.94	494.76

Note: Observation is a census tract. Specification corresponds to column (3) of Table 2. Column 1 reproduces the baseline results. Column 2 instruments the true change in CMA (i.e. including the location itself in the summation and measure employment and population in both periods instead of holding them constant at their initial values) with the baseline change in CMA measure from column 1. Column 3 instruments for the change in CMA using the average change in CMA across the IV and tram instruments constructing excluding the tract itself in the summation, while column 4 excludes all tracts closer than 1.km from the tract. In this specification, only census tracts further than 500m from a portal and a dummy for whether a census tract is further than 1km from the historical tram system is included (to capture direct effects from the tram instrument). Column 1 reports HAC standard errors as in the baseline specification. Columns 2-4 report heteroscedasticity robust standard errors. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 6: Aggregate Results Using Sufficient Statistics Approach

Panel A: VTTS Comparison			
	VTTS	GE	GE (No Ext.)
Welfare Gain (%)	1.260	2.335	1.527
90% CI	(0.740,3.106)	(0.740,5.861)	
95% CI	[0.309,3.692]	[0.475,6.936]	
As Fraction of VTTS		53.95	82.49
90% CI		(44.20,71.74)	
95% CI		[39.22,85.96]	

Panel B: Aggregate Effects		
	No Migration	Migration
Welfare	2.282	0.597
GDP	3.121	15.131
GDP Net of Costs	2.504	14.514
Population	0.000	9.514
Rents	-0.672	5.283
% of Obs GDP Growth	2.963	14.362
% of Obs Population Growth	0.000	34.886

Panel C: Incorporating Congestion		
	% Change in Welfare	% of No Congestion Welfare Change
No Congestion	3.921	100.00
Congestion	3.943	100.55
Convert TM to Car Lanes	0.028	0.641

Notes: Table shows the welfare effects from TransMilenio using the sufficient statistics approach from Proposition 1. Panel A compares the GE welfare effects with those from the first order approximation (VTTS) in proposition 2. The % change in welfare is computed as adding TransMilenio back to the counterfactual equilibrium without it. Each entry is computed by first simulating the effect of removing TransMilenio (the initial equilibrium) and then adding it back in under the different approaches. In column 1, the change in travel times accounts for the discrete choice over modes used to aggregate mode-specific travel times. In column 2, the GE effects are reported using the main reduced form elasticities (column 3 in Table 2). 90% and 95% confidence intervals are provided by bootstrapping the quantitative exercise 200 times as described in Appendix C.9. The second row reports the fraction of GE gains are explained by VTTS, with confidence intervals also reported. The non-parametric test on the bootstrap sample of whether the fraction of gains explained by VTTS is greater than 1 rejects this null (p-value 0.04). Column 3 reports GE results from a model without externalities. It computes the reduced form elasticities using the expressions derived in Appendix C.8.1, using estimates for $\theta, \alpha, \beta, \sigma$ and setting $\mu_A = \mu_U = 0$. Confidence intervals are not reported since this removes sampling variation from the 4 estimated reduced form elasticities. Panel B shows the (negative of the) value of the percentage change in each variable from removing the TransMilenio network (phases 1 through 3) from the 2016 equilibrium, under both assumptions on population mobility. The scenario with migration assumes a migration elasticity of $\rho = 3$ (see Appendix E.1 for details). The last two rows show the fraction of observed growth of population and GDP between 2000 and 2016 that can be accounted for by TransMilenio under each scenario. Bogotá's GDP increased by 105.35% (average annual growth rate of 4.6%) while population grew by 27% over the period. GDP net of costs shows the NPV increase in GDP accounting for capital costs and the NPV of operating costs as described in Appendix F.4. Note the average welfare value in Panel B differs from that in Panel A, which uses the counterfactual equilibrium without TransMilenio as the initial equilibrium for ease of comparison with the VTTS. Lastly, Panel C reports welfare results from model allowing for congestion (see Appendix E.2 for details). A congestion elasticity of 0.06 is used, the average congestion elasticity estimated for Bogotá by Duranton and Akbar (2017). The first row shows the welfare effect (the absolute value of $\bar{U}^{NoTM} / \bar{U}^{TM} - 1$) in the closed city model in this model extension, when the congestion elasticity is set to zero. This differs slightly from the baseline number since the congestion elasticity is used when calibrating the unobserved traffic matrix for the observed equilibrium, and the construction of commute times is slightly different due to the routing model of commutes. The second row shows the welfare impact of TransMilenio with congestion, and the second column shows the welfare gains as a fraction of the baseline case without congestion in row 1. The third row shows the welfare impact had TransMilenio routes been made into car lanes instead of BRT (the absolute value of $\bar{U}^{NoTM} / \bar{U}^{ReplaceTMWithRoads} - 1$).

Table 7: Commuting Elasticities

	OLS	IV-LCP&Tram	PPML	PPML	PPML
HighSkill X In Commute Cost	-0.0250** (0.0116)	-0.0295** (0.0120)	-0.0154*** (0.0028)	-0.0054 (0.0104)	-0.0253*** (0.0089)
LowSkill X In Commute Cost	-0.0278** (0.0121)	-0.0460*** (0.0150)	-0.0292*** (0.0027)	-0.0534*** (0.0109)	-0.0663*** (0.0096)
<i>N</i>	1,738	1,738	1,444	2,608	4,032
Years	1995,2015	1995,2015	2015	1995,2015	1995,2011,2015
Origin-Destination-Skill-Car Ownership FE	X	X		X	X
Destination-Skill-Year FE	X	X	X	X	X
Origin-Skill-Car Ownership-Year FE	X	X	X	X	X

Note: Outcome is the conditional commuting shares. Observation is an origin-destination-skill-car ownership-year cell. Skill corresponds to college or non-college educated workers. Only trips to work during rush hour (5-8am) by individuals aged 18-55. Columns 1 and 2 estimate OLS and IV models between 1995 and 2015. Columns 3-5 run PPML models on alternative sets of years: 2015, 1995 and 2015, and 1995, 2011 and 2015 respectively. Since the coefficient for high-skill workers is imprecise in the main specification using two years in column 4, the final column 5 pools data from 3 years. Travel times are measured according to the network in each year e.g. travel times for TransMilenio in 2011 come from the 2006 network, while those in 2015 come from the 2012 network. Standard errors are clustered at the origin-destination locality. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 8: GMM Results

Parameter	Estimate
Panel A: Firms	
μ_A	0.242* (0.128)
Panel B: Workers	
η_L	2.070** (0.805)
η_H	2.250*** (0.734)
μ_U^L	0.730** (0.372)
μ_U^H	1.002*** (0.335)

Note: Estimates are from joint GMM procedure as described in text. In Panel A, controls are locality fixed effects, log distance to CBD interacted with region fixed effects, basic tract controls (log area, log distance to the main road, log distance to a main road interacted with log distance to the CBD, dummies for each quartile of 1993 population density, 1990 employment share (employment divided by employment plus population), and 1993 college share), land market controls (share of land developed, floor area ratio, share of floorspace used for commercial purpose, and log average floorspace value in 2000) and historical controls (dummies for each quartile of population density in 1918). In Panel B, controls are locality fixed effects, basic tract controls (log area, log distance to the main road, log distance to a main road interacted with log distance to the CBD, dummies for each quartile of 1990 employment share), land market controls (share of land developed, floor area ratio, share of floorspace used for commercial purpose, and dummies for each quartile of floorspace value in 2000) and historical controls (dummies for each quartile of population density in 1918). All controls and fixed effects are interacted with group-specific dummies. Tracts closer than 500m to a TransMilenio portal are excluded. Instruments exclude the tract itself in summations, and are averages across the LCP and Tram measures as in the reduced form results. Standard errors clustered by tract obtained from 200 block-bootstrapped replications resampled at the tract-level.* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 9: Main Quantitative Results & Distributional Effects

	Average Welfare	Inequality
Panel A: Main Results		
Diff θ, η , Imperf Sub	1.007	0.546
90% CI	(0.254, 4.022)	(-0.159, 0.875)
95% CI	[0.017, 4.407]	[-0.441, 1.138]
P-value $\widehat{U}_H > \widehat{U}_L$	0.15	
Panel B: Decomposing the Role of Elasticities		
Same η, θ , Perf Sub	2.191	-0.371
Diff θ , same η , Perf Sub	2.261	-0.163
Diff θ, η , Perf Sub	2.510	-0.139
Panel C: Model Extensions		
Domestic Services	0.832	0.585
Local Home Ownership	0.854	0.596
All Renters	0.867	0.619

Note: Table shows the percentage welfare and inequality change (defined as $\widehat{U}_H/\widehat{U}_L$) from TransMilenio under models. Each entry is computed by first simulating the effect of removing TransMilenio, and reports the absolute value of the percentage welfare change from moving from the TM to no TM equilibrium. Panel A reports results from the full model, where θ_g, η_g are set to their estimated values and $\sigma = 1/0.7$ as described in the text. Confidence intervals from 200 bootstrap replications are reported (using the same procedure as described in Appendix C.9), as well as the p-value from a non-parametric test of whether the high-skill gain more than the low-skilled across these bootstraps. Panel B reports results decomposing the role of these elasticities. The first row assumes θ, η are equal across groups (set to their average value) and labor types are perfect substitutes in production. The second and third rows allow θ and η to differ across groups (set to their estimated values). Panel C shows results from model extensions to allow for employment of the low-skilled in domestic services, as well as alternative assumptions over home ownership. See Appendix E.5 and E.6 for further details.

Table 10: Policy Counterfactuals

Panel A: Alternative Networks

	% Δ Welfare	% ΔInequality	% ΔOutput
Remove Line South	-0.298	-0.060	-0.318
Remove Line North	-0.084	-0.204	-0.699
Remove Feeders	-0.942	-0.196	-1.014
Faster TM	1.355	0.698	2.790

Panel B: Land Value Capture Welfare Effects

	% Increase Relative to Baseline	
	Welfare	Output
Free Adjustment	44.04	15.78
LVC, Bands	24.47	9.17
LVC, CMA	43.82	11.95

Panel C: Land Value Capture Revenue Effects

	Closed City	Open City
LVC Band Revenue (mm)	58.62	152.77
As share of capital costs	4.04	10.54
LVC CMA Revenue (mm)	88.31	297.57
As share of capital costs	6.09	20.53


Note: Panel A shows the impact of particular network components relative to the full network using the full model. The numbers report the percentage change in each variable from moving from the full TransMilenio network to the counterfactual one. The last row reports results from making the TransMilenio faster, with an operational speed of 35km/h. Panel B shows the impacts of alternative housing supply models, using the model extension from Section E.4. I first solve for the counterfactual equilibrium without TransMilenio. I then compute the equilibrium returning to the TransMilenio network under each housing supply model, and report the percentage change in each variable as a fraction of returning to the observed network under the fixed housing supply assumption (minus one, since the change in each variable in each counterfactual scenario exceeds the value under fixed housing supply). The first row is the case with freely adjusting housing. The second row is the distance-band based land value capture (LVC) scheme, where the government sells rights to construct up to 30% new floorspace in tracts closer than 500m from stations. The third row is the CMA-based scheme where the same number of permits are issued by distributed instead by a tract's relative change in CMA as described in the text. These figures are all from the closed city model, relative comparisons are similar in the open city model. Panel C shows the government revenue earned under the land value capture policies, in levels and as a fraction of TransMilenio's construction costs. These are reported for the closed and open city model separately since the results vary by assumption. Numbers in millions of 2016 USD.

Figures

Figure 1: Change in Commuter Market Access from TransMilenio

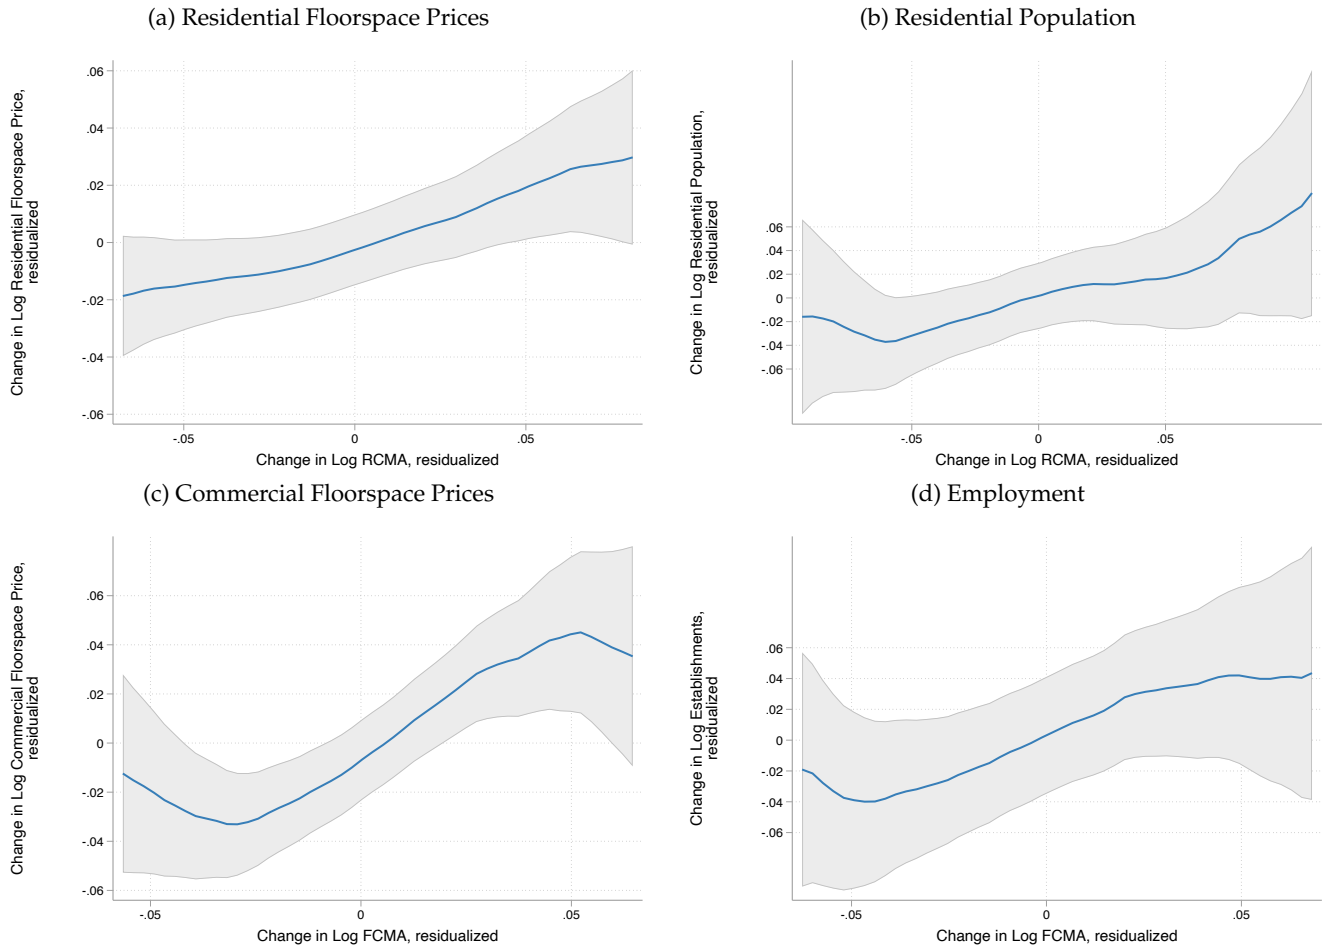
(a) Resident CMA

(b) Firm CMA




Note: Plot shows the change in CMA from the baseline specification. Population and employment are fixed at their initial level and changing only commute costs (to the full TransMilenio network as of phase 3). Tracts are grouped into vigintiles based on the change in CMA, with warmer colors indicating a larger increase in CMA. Black line shows the TransMilenio routes as of 2013. The changes in CMA are normalized to have mean zero. For the change in RCMA, the min is $-.198$, the max is $.375$, the standard deviation is 0.097 and the average range of each vigintile is $.028$. For the change in FCMA, the min is $-.147$, the max is $.246$, the standard deviation is 0.068 and the average range of each vigintile is $.020$.

Figure 2: Non-Parametric Relationship Between Outcomes and Commuter Market Access



Note: Plot shows the non-parametric relationship between outcomes and CMA. Specifications correspond to the reduced form from column (3) of Table 2. Top and bottom 2% of the change in CMA are trimmed to reduce noise at the tails and zoom in on main relationship.

Figure 3: Residential Floorspace Price Event Study



Note: See discussion in Section 4.3 for details. The year before opening is the omitted category. Heteroscedasticity and Autocorrelation Consistent (HAC) standard errors (Conley (1999)) with a 0.5km bandwidth reported.