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

Nick Tsivanidis[†]
University of California, Berkeley

May, 2022

Abstract

This paper estimates the effects of improving public transit infrastructure on city structure and welfare. It develops a sufficient statistics approach common to a class of quantitative urban models, and estimates the elasticities needed to implement the approach using data and variation from the construction of the world's largest Bus Rapid Transit system in Bogotá, Colombia. These models perform well in explaining the adjustment of economic activity to the system. The standard approach to measure the welfare gains from new infrastructure based on the value of travel time saved only accounts for 57.5% of the total welfare gain.

---

[*] A longer version of the paper that assesses the distributional impacts of the Bus Rapid Transit system can be found here. 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 helpful comments. I also thank 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 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. Any 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 and Department of Economics, University of California at Berkeley. Email: ntsivanidis@berkeley.edu. First version: November 2017.

1 Introduction

How large are the economic gains to improving public transit systems within 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),[2] measuring the benefits of these systems is challenging. Individuals may choose newly accessible home and work locations, the spatial distribution of economic activity in the city might change, and wages and house prices may respond to this reallocation. This suggests impacts that may be missed by time savings, and indirect effects throughout the city even on individuals who don't use the system. The lack of detailed intra-city data in less developed countries coinciding with the construction of large transit systems makes the task of evaluating their causal impact even more daunting.

This paper exploits uniquely detailed spatial data spanning the opening of the world's largest Bus Rapid Transit (BRT) system–TransMilenio–in Bogotá, Colombia to make three contributions to our understanding of the impact of urban transit infrastructure on cities. First, it shows that a wide class of quantitative urban models deliver a set of sufficient statistics that determine the response of economic activity and aggregate outcomes in a city to new transit infrastructure. 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. Second, it shows how changes in CMA can be measured from readily available data, and estimates the elasticities 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 with the VTTS to isolate the importance of reallocation and general equilibrium effects in valuing the gains from new transit infrastructure.

The VTTS only account for around 57.5% of the total welfare gain from the BRT. Accounting for equilibrium effects therefore matters for valuing the gains from new transit infrastructure in cities. The framework and parameter estimates from this paper provide a simple blueprint to do so.

Opened in 2000, TransMilenio is the world's most used BRT system with a daily volume of over 2.2 million trips. The system operates more like a subway than the informal bus system that preceded it: 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] I collect new sources of data covering

---

[1] McKinsey (2016) suggest 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) suggests the average subway system costs $27.81bn in 2017 dollars.

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

[3] For example, the 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).

2,800 census tracts on residence, employment, commuting patterns, and land markets spanning the system's construction.

A large literature estimates treatment effects of transit based on distance to stations, but the assumptions necessary for identification may be invalidated when there are spillovers between treatment and control units (Donaldson 2015). This is especially likely to be the case within cities where locations are interlinked through a commuting network. This paper develops a simple quantitative urban model in the spirit of recent work. The key insight is that many such models admit a log-linear reduced form in which equilibrium outcomes such as residential populations, employment and floorspace prices are log-linear functions of CMA, which summarize the network linkages across locations. The class of models includes those with endogenous firm location choice, endogenous housing supply, capital in the production function, Eaton and Kortum production, leisure in utility, individual heterogeneity in preference and productivity shocks across locations, and alternative residential and employment supply elasticities and timing assumptions. In all of these 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 may have been 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 as well as the hypothetical change that would have occurred 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 do not affect the stability of the estimates), 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 show 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 changes in surrounding unobservables. Fourth, I condition on distance to the closest TransMilenio station to assess whether effects are driven by changes in accessibility rather than other features of stations (e.g. changes in foot traffic or pollution). Fifth, I construct cost-shifting instruments to predict TransMilenio's routes based on a historical tram network and engineering estimates of the cost to build BRT on different types of land.

After showing the log-linear relationships between changes in outcomes and CMA predicted by the model are borne out in the data, the final part of the paper uses the estimated model to quantify the welfare effects of the new infrastructure. A key theoretical result that arises through an appli-

---

[4] More 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.

cation 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 account for 57.5% 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 2.21% in the baseline case where the BRT does not cause migration into Bogotá from the rest of Colombia, and 0.50% with migration. GDP per capita rose by 2.36-4.12% in these cases respectively, net of construction and operating costs. Overall, TransMilenio can account for between 2.83-12.06% 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.

I examine the robustness of these findings to alternative parameter values and a model extension that incorporates congestion. This extension builds of the framework of Allen and Arkolakis (2021), using the congestion elasticity estimated by Akbar and Duranton (2017) for Bogotá. While the data do not suggest road speeds for cars and regular buses changed on routes adjacent to the BRT relative to others, the system may have changed the overall level of speeds across the city. Allowing for congestion leads to a 30% larger gain from the system. The substitution of travelers from cars and existing buses onto the BRT frees up road space and reduces travel times for other modes. The baseline welfare effects therefore are likely to provide a lower bound of the BRT system's impact.

Two sets of counterfactuals are run to 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. While similar schemes have been used with great success in Asian cities such as Hong Kong and Tokyo, one of the main criticisms of TransMilenio was that the city experienced such a large change in transit without any adjustment of zoning laws to allow housing supply to respond. A well-targeted scheme would have increased the welfare gains from TransMilenio by around 19%, while government revenues cover 25-39% 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.

This paper contributes to several literatures. Most closely related is the body of work that 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 CBD (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 in the presence of spillovers between spatial units.[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; 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; 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). The main contribution of this paper is to show how a class of quantitative urban models admit a simple sufficient statistics approach that specifies changes in economic activity in response to new transit infrastructure in cities. It combines rich microdata with the construction of the world's largest BRT system to validate these models' key predictions and estimate the models' parameters.

Lastly, this paper relates to work in transportation economics measuring the benefits of improved transportation through the value of travel time saved (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 Bogotá and TransMilenio as well as the data. Section 3 presents a simple quantitative urban model and derives the reduced form framework and sufficient statistics approach. Section 4 estimates the model and Section 5 quantifies TransMilenio's impact. Section 6 concludes.

2 TransMilenio: The World's Most Used BRT System

Background Bogotá is the political and economic center of Colombia, accounting for 16% and 25% of the country's population and GDP respectively. In 1995 the average trip to work took 55 minutes, more than double that in US cities. The vast majority were taken by bus (73%), followed by car (17%)

---

[5] In 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 may not be externally valid in another. The CMA approach predicts different treatment effects from different transit networks based on the specific network structure.

[6] In trade and economic geography settings, firm and consumer market access often equal each other due to a balanced trade assumption (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 that the same number of people live in a location that work there, which clearly fails in the data) that delivers the BRT's very distinct shocks to resident and firm CMA shown in Figure 1.

[7] Rosenthal 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).

and walking (9%).[8] Despite its importance, public transportation in the city was highly inefficient due in large part to its industrial organization. The government allocated the administration of routes to companies called "afiliadoras" which acted as intermediaries between the government and bus companies. Afiliadoras sold slots to run their routes to bus operators. Since their profits depended only on the number of buses the result was a huge over-supply of vehicles. Low enforcement meant that up to half of the city's bus fleet operated illegally (Cracknell 2003).[9] Disregard of bus stops promoted boarding and alighting along curbs, further reducing traffic flows.

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, its first phase opening a mere 21 months later adding 42 km along Avenida Caracas and Calle 80, two arteries of the city.[10] 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 patronized system of its kind in the world (Cervero et. al. 2013).[11] Its average operational speed of 26.2kmh reported during phase one is on par with that of the New York subway (Cracknell 2003), and provided a pronounced improvement on reported bus speeds of 10kmh 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. In contrast to the informal network that preceded it, 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, as well as 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 (who live further from the CBD) while the government sets fares close to those offered by existing buses.[12]

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 these features have led to systems being built in more than 200 cities, the vast majority con-

---

[8] Bicycles and motorbikes account for the remaining 1% of commutes. For comparison, the average commute in US cities was 21 minutes in 1980 to 26 minutes in 2015.

[9] The Department of Mobility estimated the number to be more than double the amount actually required. A typical way bus companies avoided government controls was duplication of license plates and vehicle documentation.

[10] While 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). Two years prior to TransMilenio, Peñalosa implemented a "pico y placa" driving restriction which restricted cars to 3 days of peak hour weekday road based on their license plate endings (this was later extended to all day in 2009). While the main change occurred before the period of interest, my controls for locality fixed effects and distance to CBD capture potential trends in the benefit of access to public transport that vary across space induced by the policy. Although speculative, the policy did not appear to have the intended consequence of reduced car use: rising pollutants suggest increased purchases of old vehicles (Lawell et. al. 2017) and there was actually a mild increase in car ownership over the period.

[11] A map of each system component and their opening date is provided in Appendix Figure A.1, while Appendix 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.

[12] For example, in 2011 (the only year where fare information is reported in the Mobility Survey), the average bus fare is 1400 COP compared to the 1700 COP fare on TransMilenio. While the fare difference of 21.4% is non-trivial, this does not reflect the free transfers across trunk and feeder lines not offered by the existing bus network.

structed 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 level 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.[13]

Three features make TransMilenio an attractive context for empirical analysis. First, since 1980 multiple administrations worked on proposals for a subway system. These plans motivate a placebo check since the actual and hypothetical networks induce different changes in CMA across the city. Second, having identified neighborhoods towards the city's periphery to be connected with the center, final routes were chosen to a large extent by the desire to minimize construction cost. Lines were placed along wide arterial roads that were cheaper to convert and determined by the the city's historical evolution. Third, TransMilenio was was rolled out so quickly primarily to complete a portion of the system 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 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 F 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 counterfactuals, I assess the impact of TransMilenio had Bogotá pursued a such a policy.

Trip Characteristics Appendix F summarizes how TransMilenio is used. First, it is a quantitatively important mode of transit used 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, the BRT is used more for commutes to work than leisure trips when compared to other modes. TransMilenio's outsized role in commuting motivates the focus on access to jobs in this paper.

3 Using Theory to Measure the Impacts of Improving Transit in Cities

Better transit systems allow commuters to save time but may also change choices of where to live and work and have equilibrium impacts on land and labor markets throughout the city. This section develops a simple quantitative model of a city that allows for these channels yet admits a simple log-linear reduced form representation. A single measure—CMA—summarizes the effect of a city's entire transit network on any location. CMA, the elasticities of equilibrium variables to CMA, and the elasticity of commute choices to commute costs, turn out to be sufficient statistics to quantify the

---

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

impact of new transit infrastructure on economic activity in a city.

3.1 Model

Setup. I consider a simple quantitative model of a city in the spirit of Ahlfeldt et. al. (2015) and Allen et. al. (2015). There are $i \in I$ locations that differ in their exogenous amenities $\bar{u}_i$, productivities $\bar{A}_i$, residential and commercial floorspace supplies H_{Ri}, H_{Fi} and the time t_{ij} it takes to commute to any other location.[14] A continuum of workers with mass $\bar{L}$ choose where to live and work and have Cobb-Douglas preferences over a freely-traded numeraire good and housing. Commuting reduces effective labor supply at workplace so that an individual living in i and working in j receives income w_j/d_{ij}, where $d_{ij} = \exp(\kappa t_{ij})$ converts commute times into commute costs. In each location, a representative firm produces a freely traded variety under perfect competition that are aggregated by consumers in CES fashion to form the final numeraire good.

Individuals. Indirect utility across pairs of residential and employment locations (i, j) is given by

$$U_{ij}(\omega) = \frac{u_i w_j r_{Ri}^{\beta-1}}{d_{ij}} \epsilon_{ij}(\omega), \tag{1}$$

where $\epsilon_{ij}(\omega)$ is an idiosyncratic productivity for worker ω on commute (i, j), $1-\beta$ is the expenditure share on housing, and u_i is the amenity enjoyed by residents who live in i. To allow for the possibility of local spillovers, amenities depend on both exogenous location characteristics $\bar{u}_i$ and the number of residents through $u_i = \bar{u}_i L_{Ri}^{\mu_U}$. Workers choose the commute pair that maximizes their utility. Assuming these are drawn iid from a Frechet distribution with shape parameter θ yields a simple expression for the number of commuters for each live-work pair

$$L_{ij} = \bar{L}\bar{U}^{-\theta} \left(\frac{u_i w_j r_{Ri}^{\beta-1}}{d_{ij}} \right)^{\theta}, \tag{2}$$

where $\bar{U} = \gamma \left[\sum_{ij} (u_i w_j r_{Ri}^{\beta-1}/d_{ij})^{\theta} \right]^{1/\theta}$ is average utility, $\gamma = \Gamma\left(\frac{\theta-1}{\theta}\right)$ and $\Gamma(\cdot)$ is the Gamma function. The supply of residents and workers to each location can be computed by summing these flows over all destinations and origins respectively to get

$$L_{Ri} = \bar{L}\bar{U}^{-\theta} \left(u_i r_{Ri}^{\beta-1} \right)^{\theta} \Phi_{Ri} \tag{3}$$

$$L_{Fj} = \bar{L}\bar{U}^{-\theta} w_j^{\theta} \Phi_{Fj}. \tag{4}$$

The Φ_{Ri} and Φ_{Fi} terms are what I refer to as commuter market access terms. Residential commuter market access (RCMA) $\Phi_{Ri} = \sum_j (w_j/d_{ij})^{\theta}$ reflects residents' access to well-paid jobs from

---

[14] Appendix F shown housing supply was unaffected by TransMilenio, so I consider these as fixed location characteristics. This assumption is relaxed in Section 5.3. Appendix F also shows that there were no significant relative changes in car and bus speeds along routes most affected by TransMilenio, so I assume travel times are fixed in the baseline model. This is relaxed in Section 5.2.

location i. Firm commuter market access (FCMA) $\Phi_{Fj} = \sum_i (u_i r_{Ri}^{\beta-1}/d_{ij})^\theta$ reflects firms' access to workers from location j (i.e. being close to locations with high amenities or low rents). The resident supply curve (3) therefore tells us that more residents will move to locations with high amenities, low house prices, and better access to well-paid jobs through the commuting network. The labor supply curve (4) tells us that firms will attract more workers to locations with high wages and better access to workers via the commuting network.

The supply of effective labor units to a location can be computed by leveraging that, under the Frechet distribution, the average productivity of workers who have chosen (i, j) is inversely related to the share of workers choosing that pair $\bar{\epsilon}_{ij} \propto {\pi_{ij}}^{-1/\theta}$ where $\pi_{ij} = L_{ij}/\bar{L}$. Total effective labor supply is simply $\tilde{L}_{Fj} = \bar{L}\sum_i \pi_{ij}^{\frac{\theta-1}{\theta}}/d_{ij}$, which simplifies to

$$\tilde{L}_{Fj} = \bar{L}\bar{U}^{-(\theta-1)} w_j^{\theta-1} \tilde{\Phi}_{Fj} \tag{5}$$

where $\tilde{\Phi}_{Fj} = \sum_i (u_i r_{Ri}^{\beta-1})^{\theta-1} d_{ij}^{-\theta}$ is adjusted FCMA capturing access to effective units of labor.

Consumers spend a constant fraction $1-\beta$ on housing, so that residential floorspace (inverse) demand is given by

$$r_{Ri} = \frac{1-\beta}{H_{Ri}} \bar{y}_i L_{Ri}, \tag{6}$$

where $\bar{y}_i \equiv \Phi_{Ri}^{1/\theta} L_{Ri}^{-1/\theta}$ is average income of residents in i.[15]

Firms. The production side of the model assumes an Armington structure with no trade costs. In each location, a representative firm produces a differentiated variety using the Cobb-Douglas technology $Y_i = A_i \tilde{L}_{Fi}^{\alpha} H_{Fi}^{1-\alpha}$. As for amenities, I allow for the possibility of productivity externalities of the form $A_i = \bar{A}_i \tilde{L}_{Fi}^{\mu_A}$.[16] Solving firms' profit maximization problem delivers labor demand

$$\tilde{L}_{Fi} = \alpha w_i^{\alpha(1-\sigma)-1} A_i^{\sigma-1} r_{Fi}^{(1-\sigma)(1-\alpha)} E \tag{7}$$

where $E = \sum_i \bar{y}_i L_{Ri}$ is aggregate expenditure and σ is the elasticity of demand across varieties. Firm (inverse) demand for commercial floorspace is given by

$$r_{Fi} = \left((1-\alpha) \frac{A_i^{\sigma-1} w_i^{-\alpha(\sigma-1)} P^{\sigma-1} E}{H_{Fi}} \right)^{\frac{1}{1+(\sigma-1)(1-\alpha)}} \tag{8}$$

Equilibrium. Given model parameters $\{\alpha, \beta, \sigma, \theta, \kappa, \mu_U, \mu_A\}$ and location characteristics $\{H_{Ri}, H_{Fi}, t_{ij}, \bar{u}_i, \bar{A}_i\}$, an equilibrium of the model is a vector $\{L_{Ri}, \tilde{L}_{Fj}, w_j, r_{Ri}, r_{Fj}, \bar{U}\}$ such that (i) the supply of residents

---

[15] See Appendix C.6 for a derivation. The model with separate residential and employment location decisions covered in Appendix C.4 has the more familiar form $\bar{y}_i \equiv \Phi_{Ri}^{1/\theta}$.

[16] Given evidence on highly localized spatial spillovers (Rossi-Hansberg et. al. 2010; Ahlfeldt et. al. 2015), I do not allow for spillovers across locations given the size of census tracts. Previous versions of the paper show how the regression framework in that model still holds but outcomes depend both on a location's own CMA and those nearby.

and labor is consistent with worker optimality (3) and (5), (ii) the demand for labor is consistent with firm optimality (7), (iii) demand for floorspace is consistent with firm and worker optimal and equals supply (6) and (8) and (iv) the population of the city $\bar{L}$ is fixed, and welfare $\bar{U}$ is given by $\bar{U} = \gamma \left[\sum_{ij} (u_i w_j r_{Ri}^{\beta-1} / d_{ij})^\theta \right]^{1/\theta}$.[17]

3.2 Quantifying the Impacts of Transit Infrastructure

The following proposition shows how the model and related extensions admit a simple reduced form and sufficient statistics approach to quantify the impacts of changes in transit infrastructure. The proof of this and subsequent propositions are provided in Appendix C.

Proposition 1. *Consider a change in commute costs from* $\mathbf{d}$ *to* $\mathbf{d}'$*, and let* $\hat{x} \equiv x'/x$ *denote relative changes in a variable between the pre- and post-period. Then*

Part 1: Reduced Form. *The model yields a reduced form where endogenous variables can be written as log-linear functions of CMA*

$$\ln \hat{\mathbf{y}}_i = \boldsymbol{\beta}_R \ln \hat{\Phi}_{Ri} + \boldsymbol{\beta}_F \ln \hat{\Phi}_{Fi} + \mathbf{e}_i$$

where $\mathbf{y}_i = [L_{Ri}, r_{Ri}, r_{Fi}, L_{Fi}]$ *and* $\mathbf{e}_i$ *is a vector of structural residuals.* $\boldsymbol{\beta}_F$ *and* $\boldsymbol{\beta}_R$ *have zero elements in the first and last two entries respectively, so this is a system of 4 univariate regressions yielding 4 coefficients* $\beta_{L_R}, \beta_{r_R}, \beta_{r_F}, \beta_{L_F}$*. Unique (to-scale) values of the CMA terms* Φ_{Ri}, Φ_{Fi} *can be computed given data* $\{L_{Ri}, L_{Fi}, d_{ij}\}$ *and the commuting elasticity* θ*.*

Part 2: Relative Impacts of Transit Infrastructure*. Assuming that exogenous, location-specific characteristics are unchanged by the infrastructure, relative changes in endogenous variables* $\hat{\hat{\mathbf{y}}}_i \equiv \hat{\mathbf{y}}_i / \left(\prod_r \hat{\mathbf{y}}_i\right)^{1/I}$ *can be computed using (i) estimates of* $\beta_{L_R}, \beta_{r_R}, \beta_{r_F}, \beta_{L_F}, \theta$*, (ii) data on the initial distribution of economic activity* $\{L_{Ri}, L_{Fi}, d_{ij}\}$ *and (iii) data on the change in commute costs* $\{\hat{d}_{ij}\}$*.*

Part 3: Level Impacts of Transit Infrastructure*. Level changes in endogenous variables* $\hat{\mathbf{y}}_i$ *and endogenous constants* $\hat{\bar{L}}, \hat{\bar{U}}$ *can be computed from the relative changes obtained in part 2 with (i) an assumption on population mobility between the city and the rest of the country, and (ii) values for* σ, β*.*

Part 4: Isomorphisms*. Parts 1 and 2 apply to a more general class of models which feature (i) a gravity equation for commute flows and (ii) an equilibrium that can be written as a system of* K *equations in* K *endogenous variables* $\{y_{1i}, \ldots, y_{ki}\}_{i=1}^{I}$ *of the form*

$$\prod_{k=1}^{K} y_{ki}^{\alpha_{kh}} = \lambda_h \Phi_{Ri}^{b_h^R} \Phi_{Fi}^{b_h^F} e_{ih} \quad \text{for } h = 1, \ldots, K.$$

These models will yield the same counterfactual changes in outcomes (relative to city-wide averages) as those from the baseline model, given estimates of β_R, β_F, θ*. This class includes models with endogenous firm location choice, Eaton and Kortum production, capital in the production function, endogenous housing supply, leisure,*

---

[17] Existence of the equilibrium and conditions for uniqueness are established in Proposition 3 in the appendix. Alternative assumptions over population mobility between Bogotá and the rest of the country are covered in Proposition 1.

preference rather than productivity shocks, and alternative residential and employment supply elasticities and timing assumptions. However, the overall level of changes and changes in endogenous constants will depend on (a subset of) the particular structural parameters of the model $\{\{\alpha_{kh}\}_k, b_h^R, b_h^F\}_h$*, and are not determined by the reduced form elasticities alone.*

The implications of these results are now discussed in turn.

Reduced Form Representation. The first part of Proposition 1 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 between two periods depends only on the change in CMA as well as a structural residual that reflects changing location fundamentals (productivities, amenities and floorspace supplies).[18] This system reduces to a system of 4 univariate regressions, where residential outcomes depend on RCMA and commercial outcomes depend on FCMA

These CMA terms can be easily recovered using data on residential populations, employment, commute costs and the commuting elasticity θ. This ensures estimation of the reduced form is straightforward, even if CMA is not directly observed in the data. The proof of Proposition 1 shows that the CMA terms are the unique to-scale solution to the system

$$\Phi_{Ri} = \sum_j d_{ij}^{-\theta} \frac{L_{Fj}}{\Phi_{Fj}} \tag{9}$$

$$\Phi_{Fj} = \sum_i d_{ij}^{-\theta} \frac{L_{Ri}}{\Phi_{Ri}}. \tag{10}$$

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 wage 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 population, particularly so when these other locations lack access to jobs (lowering the wage individuals are willing to work there for).

To visualize the change in CMA and how it differs from distance-based measures of treatment effects standard in the literature, Figure 1 plots the distribution of changes in commuter access across the city induced by the construction of the first two phases of the system.[19] The system increases access to jobs much more for tracts in the outskirts of the city, which were far from the high-employment densities towards the center. Firms' access to workers rose more in the center, since these locations were best positioned to take advantage of increased labor supply along all spokes of the network.[20] Moreover, this shows how all locations in the city are treated by the BRT system due to the realloca-

---

[18] The contents of the residual and reduced form parameters are outlined in Appendix C.5. The residual contains changes in unobserved amenities and residential floorspace for residential outcomes, and changes in unobserved productivities and commercial floorspace for commercial outcomes.

[19] The figure plots the change in CMA induced by holding population and employment fixed at their initial level in 1993 and 1990 respectively (from the population and economic census) and changing only commute costs to isolate graphically the change due only to TransMilenio (i.e. the main measure discussed in Section 4.3).

[20] FCMA increases toward the center-North due to the high density of (low-skill) workers in the South.

tion of employment and residential population across the city.

Counterfactual Impacts of Transit Infrastructure. Part 2 of Proposition 1 shows that relative changes in endogenous variables across the city in response to a change in commute costs can be computed using data on the initial distribution L_{Ri}, L_{Fi}, d_{ij}, the change in commute costs $\hat{d}_{ij}$, the commuting elasticity θ, and the reduced form parameters $\beta_{L_R}, \beta_{r_R}, \beta_{r_F}, \beta_{L_F}$. In other words, these data and parameters are sufficient statistics for the change in economic activity across the city in response to changes in transit infrastructure. As shown in the proof, the elasticities and the change in CMA are the sufficient statistics; the data on initial economic activity and changes in commute costs are necessary to compute the change in CMA.

Part 3 shows that computing both the level change in endogenous variables as well as the change in equilibrium constants requires slightly more structure. These require an assumption on population mobility into the city from the rest of the country, and 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 to external values or aggregate moments.

Part 4 shows that some of these results apply more generally to a wider class of models which feature a gravity equation for commute flows and a log-linear equilibrium representation. Despite having different underlying structural parameters, these models yield the same log-linear reduced form. Since part 2 requires only values of these reduced form elasticities to compute relative changes in activity across the city in response to changes in the transit network, they yield the same (relative) counterfactual impacts as the baseline model. This result is particularly useful because the researcher does not need to take a stand on which particular modeling assumption is true; each will yield the same counterfactual impact on relative outcomes as the baseline model conditional on the reduced form estimates $\boldsymbol{\beta}_R, \boldsymbol{\beta}_F$. Where the modeling assumptions do come into play is in determining the overall level of changes and aggregate effects (such as welfare). As the example in part 3 shows, this depends on the underlying structural parameters of the model. However if the researcher is ready to take a stand on the value of those parameters in their model, then these aggregate impacts can be computed using the procedure shown in the proof of part 3 and the values of the particular structural parameters of that model.

4 Estimation

4.1 Data Description

This section summarizes the data used in the analysis, with further details in Appendix E.

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 Census 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. In 2015 the data 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. Appendix Figure A.4 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 which allows 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 opening.

Housing market data between 2000 and 2018 comes 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 2 million properties (Ruiz and Vallejo 2010).[21] It reports the use, floorspace and land area, 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, Appendix Figure A.5 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 member was asked to complete a travel diary for the previous day. The survey reports the de-

---

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

mographic information of each traveller 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, mode, as well as origin and destination UPZ.

Employment data at the worker level 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 each year. They report individual and household characteristics, as well as details on employment such as income, hours worked and industry of occupation across primary and secondary jobs. I use versions of these datasets with the block of each household reported.

Commute times between more than 7.8mm pairs of census tracts by each mode are computed in ArcMap. I obtain the shape of each mode's network by combining spatial datasets provided by the city. To construct the time to traverse each edge of the network, I assign speeds in order to match both reported values in the literature as well as the distribution of commute times observed in the Mobility Surveys. Appendix Figure A.6 shows the computed times correlate well with observed door-to-door times from these surveys.

4.2 Estimation Part I: Commuting Parameters θ, κ

Identifying κ. To identify κ, I extend the model to incorporate multiple travel modes.[22] After choosing their commute, workers choose which transit mode to use. There are two nests of transit modes available. $\mathcal{B}_{Pub} \equiv \{\text{Walk}, \text{Bus}, \text{TransMilenio}\}$ is the nest of public modes while $\mathcal{B}_{Priv} \equiv \{\text{Car}\}$ is the nest of private modes. Workers become car owners according to a Bernoulli distribution with parameter ρ_{car} . Letting $a \in \{0,1\}$ index whether or not an individual owns a car, they can choose to commute using modes $m \in \mathcal{M}_a$ available to them with $\mathcal{M}_0 = \mathcal{B}_{Pub}$ and $\mathcal{M}_1 = \mathcal{B}_{Pub} \cup \mathcal{B}_{Priv}$. Individuals have idiosyncratic preferences across modes such that the realized commute cost for individual ω is given by $d_{ijm}(\omega) = \exp\left(\kappa t_{ijm} + \upsilon_{ijm}(\omega)\right)$. The commuter's mode choice problem conditional on having chosen commute ij and drawing car ownership status a is simply to choose the mode available to them that minimizes commute costs $\min_{m \in \mathcal{M}_a} \{d_{ijm}(\omega)\}$.

Following precedent in the transportation literature (e.g. McFadden 1974), I assume the $\upsilon_{ijm}(\omega)$ are drawn from a generalized extreme value distribution

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

This distribution allows for correlation of preference shocks within nests, with $\lambda_k \to 0$ being the case of perfect correlation. Note that $\lambda_{Priv} = 1$ by virtue of there being only one private modes. $\tilde{b}_m$ capture attributes other than travel time that influence mode choices.

The parameters of the mode choice model are κ, λ, b_m. These are estimated via maximum likelihood using standard expressions for choice shares in the nested logit model. The data comes from

---

[22] Full derivations and additional estimation details are provided in Appendix C.7.

the 2015 Mobility Survey when all modes are available. κ is identified from the sensitivity of mode choices within commute (i, j) 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 presented in Table 1. The estimate of $\kappa = 0.013$ is very close to that of 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 as well as the inconvenience of having to walk between stations and final origins and destinations.

With these estimated parameters, I take expectations over the mode choice problem to derive the structure of the utility in (1) of the baseline model. As shown in the appendix, expected utility prior to realizing the mode-specific preference shocks is then given by (1), only with t_{ij} replaced by expected commute times $\bar{t}_{ij}$. These are given by

$$\begin{aligned} \bar{t}_{ij} &= -\frac{1}{\kappa} \ln \left[\rho_{car} \exp\left(-\kappa \bar{t}_{ij1}\right) + \left(1 - \rho_{car}\right) \exp\left(-\kappa \bar{t}_{ij0}\right)\right] \\ \text{where } \bar{t}_{ij0} &= -\frac{\lambda}{\kappa} \ln \sum_{m \in \mathcal{B}_{\text{Public}}} \exp\left(b_m - \frac{\kappa}{\lambda} t_{ijm}\right) \\ \bar{t}_{ij1} &= -\frac{1}{\kappa} \ln\left(\exp(b_{car} - \kappa t_{ij\text{Car}}) + \exp\left(\kappa \bar{t}_{ij0}\right)\right), \end{aligned}$$

where $b_m \equiv -\tilde{b}_m / \lambda_{k(m)}$. I set $\rho_{car} = 0.181$ equal to the share of car owners in 2015. I explore the sensitivity of my results to alternative ways of aggregating mode-specific times in robustness exercises.

Identifying θ. Taking logs and first differences of the expression for commute flows (2) yields a gravity equation relating the change in commute flows to changes in commute times

$$\ln L_{ijt} = \alpha_{ij} + \gamma_{it} + \delta_{jt} - \theta \kappa t_{ijt} + \varepsilon_{ijt},$$

where α_{ij}, γ_{it} and δ_{jt} are origin-destination, origin-year and destination-year fixed effects. Given the estimate of κ from the previous step, θ is identified from the sensitivity of changes in commute flows due to changes in commute times. While other estimation approaches typically leverage cross-sectional variation, I use the change in commute times induced by TransMilenio to difference out time-invariant characteristics potentially correlated with commute times. The identification assumption is that changes in commute times due to TransMilenio are uncorrelated with unobserved pair-specific factors that also determine growth in commute flows. Changes in origin- or destination-specific unobservables—such as amenities and productivities—are absorbed in the fixed effects.[23]

The results are presented in Table 2. Given the presence of zeros in the commuting data, the preferred specification estimates the relationship using PPML. The estimate in column 2 which includes

---

[23] An IV estimate for θ using approach from Section 4.3 is presented in Table A.4 and is used in robustness exercises. Repeating column 2 (my preferred estimate) using data from 2015 alone leads to a 30% smaller estimate of the commute elasticity (in absolute value), suggesting a meaningful impact of the ability to difference out pairwise unobservables.

controls for route observables interacted with year fixed effects implies a value of $\theta = 3.398$. This is similar to existing estimates (e.g. Monte et. al. 2018; Heblich et. al. 2020; Severen 2021). This will be used as the baseline value in what follows, but will be varied in robustness checks. Columns 3 and 4 repeat the exercise using OLS (with some pairs with zero commute flows dropped as a result), leading to somewhat noisier but similar point estimates. The last two columns weight each observation by the number of commuters in each origin destination cell which sharpens up the OLS estimates.[24]

4.3 Estimation II: Reduced Form Elasticities $\beta_{L_R}, \beta_{r_R}, \beta_{r_F}, \beta_{L_F}$

Part 1 of Proposition 1 implies the equilibrium can be written

$$\ln \hat{\mathbf{y}}_{Ri} = \boldsymbol{\beta}_R \ln \hat{\Phi}_{Ri} + e_{Ri} \tag{11}$$

$$\ln \hat{\mathbf{y}}_{Fi} = \boldsymbol{\beta}_F \ln \hat{\Phi}_{Fi} + e_{Fi}, \tag{12}$$

where $\mathbf{y}_{Ri} = [L_{Ri}, r_{Ri}]$ and $\mathbf{y}_{Fi} = [L_{Fi}, r_{Fi}]$ are residential and commercial outcomes. The coefficients $\boldsymbol{\beta}_R = [\beta_{L_R}, \beta_{r_R}]$ and $\boldsymbol{\beta}_F = [\beta_{L_F}, \beta_{r_F}]$ are the reduced form elasticities needed to conduct quantitative analysis of transit infrastructure construction.

As described in the discussion of Proposition 1, the structural residual contains changes in amenities and residential floorspace supplies (for residential outcomes) and productivities and commercial floorspace supplies (for commercial outcomes). This highlights the challenge in identifying the model's parameters: estimates would be biased if Bogotá's government chose routes in a way that increased accessibility in neighborhoods with differential trends in these characteristics. For example, the government may have wanted to support growing neighborhoods or to stimulate lagging ones by improving access. 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 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 effects are driven by changes in accessibility rather than other features of stations (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 as well as the change induced by these (hypothetical) planned networks. The coefficients on the planned CMA variables can be interpreted as a placebo check that planned but not-built 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 highlighted by Borusyak and Hull (2021). Fifth I exploit TransMilenio's staggered rollout across three phases through event studies and falsification

---

[24] Standard errors are clustered by origin-destination pair. Appendix Table A.3 provides results (i) clustered by origin-year and destination-year and (ii) two-way clustered by origin and destination.

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 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 therefore construct versions of the change in CMA by solving (9) and (10) 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 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.

Main Specification. Table 3 presents the baseline results. Each observation is a census tract. Each entry corresponds to the coefficient from a regression of the change in each outcome on the change in RCMA or FCMA. 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.[25] Establishment regressions are weighted by the share of establishments in 2000 in tract to increase precision.[26] Since some establishment results are noisy, I include as an additional outcome the share of floorspace used for commercial purposes to provide evidence on TransMilenio's impact on the reallocation of employment.[27] Standard errors are clustered by census tract, although robustness checks adjust standard errors for spatial correlation in the data as in Conley (1999).

The first specification in column (1) includes controls for locality fixed effects, basic tract characteristics, and log distance to CBD interacted with region dummies.[28] 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

---

[25] In 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 as the land market regressions. This is preferred to the census employment data since it gives employment 10 years to respond to TransMilenio. Table 5 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.

[26] Unweighted regressions are run in robustness checks.

[27] In the baseline model this should not change. An extension of the model in Appendix D.3 allows for floorspace use shares to respond to TransMilenio (with total floorspace supply held constant, as observed in the data). However, the model does not retain an exact log-linear reduced form, but one with heterogeneous elasticities since each location's initial floorspace use share enters its elasticity of economic activity to CMA. The model with endogenous housing supply in Appendix C.4 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.

[28] The North of Bogotá tends to contain richer and more educated census tracts than the West and South, so this interaction allows for differential growth further away from the city center across these regions.

log establishments in the final row, whose coefficient almost halves when the full set of controls are added. I consider column (3) to be the baseline specification continued in later tables, as this includes the full set of controls.

Column (4) excludes tracts which 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, and so the location of these portals may have been endogenous to underlying trends in economic activity at these locations. The coefficients remain largely stable in this subsample of tracts, suggesting endogeneity in the locations directly targeted by TransMilenio is not significantly affecting the results.

Column (5) uses the change in CMA to locations further than 1.5km away from a tract. Network additions so far from a census tract are less likely to linked to local trends in unobservables, thus addressing the endogeneity of network placement. The results remain robust and for the most part stable. Column (6) uses an alternative method to aggregate commute times where individuals take the quickest mode of public transit available. While this changes some of the point estimates and improves precision, the results are qualitatively similar.

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

One way to interpret these coefficients is to examine the implied values of the agglomeration and amenity externality they are consistent with. Appendix C.5 shows that $\mu_A = \frac{\sigma}{\sigma-1} \Big/ \frac{\beta_{L_F}}{\beta_{r_F}} - \alpha$ and $\mu_U = 1 - \beta \Big/ \frac{\beta_{L_R}}{\beta_{r_R}}$, so that these externalities can be expressed as functions of (i) the reduced form elasticities and (ii) values of α, β, σ. The implied agglomeration externality using the values of α, β, σ calibrated in Section 5.1 is $\mu_A = 0.112$. This parameter is the subject of a large empirical literature and therefore easier to benchmark. My estimate lies slightly above the 0.03-0.08 range in the survey by Rosenthal and Strange (2004) and the value of 0.07 Ahlfeldt et. al. (2015) estimate in Berlin, but within the bounds established in experimental approaches in the US with estimates as high as 0.12 and 0.2 (Greenstone et. al. 2010; Kline and Moretti 2014). The amenity externality μ_U is less well-studied. I find $\mu_U = 0.626$, around 4 times larger than that of 0.15 which Ahlfeldt et. al. (2015) estimate in Berlin. The context of a developing country city may explain the difference.[29]

Visualizing the Relationship. Figure 2 plots the non-parametric relationship between (residual) growth in outcomes and (residual) changes in CMA. The relationship appears approximately log-linear for each outcome, as predicted by the model. This suggests the model performs well in fitting the heterogeneous effects observed in the data: tracts that experience large improvements in market access report large changes in outcomes.

Hypothetical Changes in CMA from Historical Network Plans. The location of the TransMilenio network was not a random decision. The government may have located the network to support or

---

[29] For example, a lack of centrally provided public goods may incentivize more density if private actors find it more profitable to provide services (such as transport, waste removal or energy provision) to dense areas.

spur existing trends in economic activity in nearby neighborhoods. To provide additional evidence of TransMilenio's causal impact, I leverage four distinct historical plans for the TransMilenio network digitized from planning documents.

Since 1980 multiple administrations 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 suggestion of the Japanese International Cooperation Agency (JICA) to build a BRT, given the cost of the subway would be "ten times higher than the alternative of articulated buses".[30] I obtained and digitized the maps for these four planned networks, shown in Appendix Figure A.3.[31] I then solve for the predicted change in CMA had TransMilenio been been built along each planned network, and compute the average change in log RCMA and FCMA across all four plans.

I extend my baseline specification (column 3 of Table 3) 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 the selection of routes based on trends in unobservables in their neighborhoods, there should be no impact of these planned but not-built networks over and above the effect of the BRT. The second interpretation is that this controls for the omitted variable bias that can arise from a location's non-random exposure to transport infrastructure 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 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 4 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 ranges from 0.18 to 0.96 with an average of 0.5 across the five outcome variables. The coefficient on the planned CMA variable is statistically insignificant from zero in all specifications. These results suggest both that the observed impacts of TransMilenio are unlikely due to pre-existing trends in neighborhoods selected by city planners, and that the existing set of controls do a sufficient job in controlling for any omitted variable bias that ex ante may arise from non-random exposure to the network.

Staggered Station Openings. TransMilenio was opened across three phases during the 2000s and

---

[30] See https://www.transmilenio.gov.co/publicaciones/146028/historia-de-transmilenio/.

[31] See https://www.metrodebogota.gov.co/sites/default/files/documentos/Producto%2015.%20Tomo%201.%20Formulación%20y%20 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.

2010s.[32] In this section, I run a set of falsification checks to test for changes in outcomes prior to the 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 phase 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 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 (to identify $\beta^{Current}$) and (ii) phase 3 of the system open by 2013 (to identify β^{Future}).[33] While prices may experience some anticipation effects, the 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 2005 and 1993 census. The right hand side variables include CMA growth due to phase 1 (open by 2003, with most stations opening by 2001), and 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.[34]

In Table 5, 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 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 due to future lines. While insignificant, these estimates of β^{Future} can be noisy. Panel B finds similar patterns for commercial land market outcomes.

---

[32] Phase 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 2012 representing 84% of the openings

[33] Explicitly, for land market outcomes $\Delta_{t,t-\ell} \ln \Phi_i$ is the change in CMA going from the pre-TM network to the phase 2 network holding residential populations and employment fixed at their initial values. $\Delta_{t+k,t} \ln \Phi_i$ is the change in CMA going from the phase 2 to phase 3 holding holding residential populations and employment fixed at their initial values.

[34] The 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. This motivates the focus on the CCB data in the rest of the analysis.

Panel C examines the impact on total and formal employment from the economic census.[35] In the odd columns which regress on realized changes in CMA, I document positive but noisy coefficients (p-values of 0.114 and 0.098). While these estimates are larger than the baseline estimates using the CCB data, which is surprising given the shorter time frame for employment to respond, the different is statistically insignificant given the imprecision of the estimates. The even 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 $\delta_t' X_i$ are a set of controls that are allowed to have time-varying impacts on house prices. 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 these earlier lines have 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 price to the line openings. It is not clear ex ante 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, there is a mild decrease in house prices in tracts which experience a larger growth in accessibility due to phase 3 in the two or three years prior to opening, consistent with this possibility. The year before the lines open, the responsiveness of prices to CMA jumps around 0.4 log points. This is potentially due to anticipation effects as the opening of the third phase became certain.[36] This effect is stable until 2 years after opening, after which the elasticity rises 1 log point over the period until 6 years after opening.

While the difference between 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 E, part of the annual change in prices in the cadastral database is based on inflating prior years' values. Primary data is collected by the cadastral office to fully reassess properties—based on collection of information on properties for sale, making offers to elicit true sales values, and in person visit by professional assessors—but this happens fairly infrequently (around 3 times over the period in question). This motivates the focus on long-run impacts in the rest of the paper.

---

[35] Formal employment is defined as employment in establishments with 5 or more workers.

[36] Corruption 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.

Instrumental Variables to Predict TransMilenio's Placement. As a last approach to address the endogenous placement of the TransMilenio network, I construct two cost-shifting instruments for TransMilenio routes. These in turn imply two instruments for the change in CMA.[37] 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 will be 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 opened in 1884, which was last extended in 1921 and stopped operating in 1951. I extend the 1921 lines to the edge of the city in present day, 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 last well after its construction, which I capture by including historical controls.[38] 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 worry 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 6 presents the results. Column 1 reproduces the baseline results for reference. Column 2 instruments 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 (fixing residence and employment at their initial levels and excluding the location itself in the summation). Column 3 instruments the realized change in CMA using the average change across the tram and least cost path instruments. The coefficients are mostly stable across these specifications, the exception being residential floorspace prices which roughly double when moving to the final column. This supports the impression from 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 quantitative analysis but explore robustness of the results to using the elasticities from the other columns in this table.

---

[37] Additional details can be found in Appendix E.3. To compute the instruments, I first calculate the commute times had the system been built along each instrument. Plugging these into (9) and (10) 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 in the initial period without the system. As in column (5) of Table 3, I exclude tracts within 1.5km of a tract when summing to compute the instrument to increase plausibility of the exclusion restriction. Historical and least cost instruments are often used in the literature (Baum-Snow 2007; Duranton and Turner 2012; Faber 2014, Alder 2019).

[38] These include 1918 population and distance to main roads in 1933. The extension of the tram lines to the city's edge should also reduce concerns over direct effects on outcome growth, since much of the instrument was never built.

Robustness. Appendix Table A.1 assesses the robustness of these results to a number of alternative specifications. First, I use alternative ways to aggregate mode-specific commute times and alternative travel speeds on each mode (columns 2 to 4). Second, I vary the commute elasticity θ to 1.5 and 0.5 times its estimated value (columns 4 and 5). Third, I consider only tracts within 3km a TransMilenio station to ensure the results are not driven by outliers at implausible distances from the network (column 7). Fourth, I use Conley (1999) HAC standard errors (compared to the baseline estimates which cluster by census tract) to allow for arbitrary spatial correlation of errors across tracts within 500m of each other (column 8). Fifth, I exclude tracts within 1km of a portal (compared to the 500m exclusion in Table 3) to further ensure the results are not driven by the targeting of these neighborhoods (column 9). Sixth, I control for distance to a tract's closest TransMilenio station interacted with distance to the CBD (column 10). This assesses whether the CMA effect is simply due to heterogeneity of the distance effect at different distances from the CBD (a possibility given the trends in Figure 1). Reassuringly, the results are robust to this, highlighting how the key source of identifying variation is local changes in RCMA within localities.

Seventh, I run an unweighted regression for the change in establishments which is weighted by the initial share of establishments in a tract in the main results (Appendix Table A.2). The unweighted results are significant as controls are added, but become noisy and insignificant in the full specification in column 3 (p-value of 0.2). I use the weighted regressions in the main results for two reasons. First, we might expect noise in the CCB data which is a database of establishments registered with the city's chamber of commerce rather than a census. Weighting by initial shares places more weights on tracts where establishment growth is more precisely estimated. Second, I document sharp positive impacts of CMA on the share of floorspace used for commercial purposes (another measure of the changing allocation of real production activity). Taken together, these suggest employment is indeed responding to TransMilenio.[39]

Lastly, Appendix F establishes a positive impact of TransMilenio on labor income, although around half the effect is explained by changes in educational composition of workers. While there is clearly a change in sorting patterns in response to TransMilenio, these results nevertheless validate the notion of transit infrastructure improving access to well-paid jobs emphasized in the model.[40]

5 The Welfare Effects of TransMilenio

5.1 First Order vs General Equilibrium Welfare Effects

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), in which its benefits are given by minutes saved times the value of time. The following proposition shows that under certain conditions, this is precisely the first order welfare impact from a change in infrastructure in the full general equilibrium model.

---

[39] A final robustness check in Appendix F.5 tests the prediction of the baseline model that residential outcomes depend only on RCMA and vice verse for commercial outcomes.

[40] The longer working paper version provides empirical evidence on changes in sorting by education group to the system, and uses a richer model to evaluate its distributional consequences.

Proposition 2. *In a version of the baseline model with (i) no amenity or productivity spillovers, (ii) preference shocks over residential locations, (iii) workers owning an equal share of all floorspace and (iv) a labor income tax* $1/(1+\theta)$ *redistributed lump sum, the elasticity of welfare to a change in commute costs is*

$$d \ln \bar{U} = -\alpha\beta\kappa \sum_{ij} \frac{w_{ij} L_{ij}}{\sum_{rs} w_{rs} L_{rs}} dt_{ij}, \tag{13}$$

where w_{ij} *is average labor income of commuters along pair* (i, j)*.*

The proof of the proposition first establishes that under these conditions the equilibrium is efficient. An application of the envelope theorem then shows that—to a first order—only the time savings from new infrastructure matter for welfare. This is simply proportional to a labor income-weighted average of the commute time reductions, scaled by κ and $\alpha\beta$. The former converts commute times to commute costs, while the latter reflects that a share of the gains go to floorspace owners rather than directly to workers.[41] Lastly, as explained in the proof of the proposition, technical reasons require the restrictions (ii)-(iv) to be imposed to derive this result. However, simulations of small shocks in the model from Section 3.1 with only condition (i) imposed confirm this expression correctly captures the first order welfare effects in that model as well.

I now quantify the welfare gains from TransMilenio using this approximation and the full general equilibrium model. This decomposes how much accrue through the time savings approach versus other equilibrium channels in the model. Doing so requires values for α, β and σ. I estimate the $1 - \alpha = 0.206$ by computing the share of floorspace in total costs across establishments in each one digit non-agricultural industry, and averaging these by the sectoral employment shares in Bogotá.[42] 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.

Table 7 shows the results. It first uses the model to simulate what Bogotá would look like in 2018 without TransMilenio.[43] It then computes the welfare change from adding in TransMilenio under the different approaches. The first entry shows the gains from TransMilenio under the first order approximation or VTTS approach from Proposition 2. This delivers a large increase in welfare of 1.3%, accruing solely through TransMilenio's higher speed relative to other modes.[44]

The second entry shows the welfare gains using the full model from Proposition 1 and the estimated elasticities. These deliver a much larger gain of 2.26%. The VTTS therefore only account for 57.5% of the total welfare gains, yielding one of the paper's central results that equilibrium effects

---

[41] While these gains ultimately make their way back to workers who own the housing stock, these equilibrium price effects do not matter to a first order.

[42] The 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.

[43] I 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 is taken to include phases 1, 2 and 3 of the system. There may be multiple equilibria in the presence of spillovers as shown in Proposition 3. 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.

[44] These figures can also be interpreted as the equivalent variation (as a fraction of income) to keep utility constant across the observed and counterfactual equilibria, at the counterfactual equilibria prices.

matter for valuing the gains from new transit infrastructure in cities. Confidence intervals for this main welfare effect are reported from a bootstrap procedure that accounts for the uncertainty in the model's parameter estimates (see Appendix C.8 for details).

The difference between the equilibrium and 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, for example). The final entry of the table shows that when externalities in productivity and amenities are turned off, the VTTS explain a larger portion of the equilibrium effects. The size of the shock explains about one third of the 43% gap between the VTTS and general equilibrium welfare effect, with the externalities accounting for the remaining two thirds.

5.2 Main Quantitative Results

Main Results. Table 8 presents the effects of TransMilenio on GDP, aggregate rents and welfare.[45] The first column presents the closed city results from the model developed above. The second column presents results from an extension outlined in Appendix D.1 which allows for an upward sloping supply of migrants into the city from the rest of the country. I use a value of $\rho = 3$ for the migration elasticity close to that of Bryan and Morten (2019), which is varied in robustness checks.

Panel A shows 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% with a slight fall in the overall level of floorspace prices.[46] When migration is allowed, the welfare gain falls to 0.5% since the increase in population of 7.97% bid up floorspace values by 4.33%. GDP of Bogotá rises by 12.71%, but this is mostly due to population growth: GDP per capita rises by 4.74%.

How much of the observed growth in Bogotá's population and GDP growth over the period does this account for? This is shown in Panel B: TransMilenio can account for between 2.83-12.06% of GDP growth in Bogotá from 2000 to 2016, and up to 29.24% of observed population growth. This suggests TransMilenio's effects are quantitatively important, but not implausibly large.

Costs vs Benefits. How did the output gains from TransMilenio compare with its costs? Panel C of Table 8 provides a breakdown of the system's costs and benefits (see Appendix E.4 for details on cost calculations). Even using the most conservative estimate in column (1), the net present value of the net increase on GDP was about \$33bn, or a net increase of 2.36% in the steady-state level of GDP. This suggests the system was a highly profitable investment for the city.

---

[45] The table reports the absolute value of the percentage change in each variable under the counterfactual without the TransMilenio network, i.e. $100 \times (X_{2015}^{NoTM}/X_{2015} - 1)$ for any variable X_{2015}. Numbers may therefore differ from the first order approaches in Table 7 which inverts the ordering by using the equilibrium without TransMilenio as the base.

[46] It 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. I find the slight negative effect noted here is due to the low residential floorspace price elasticity relative to the residential population elasticity. This implies a large value of the residential externality $\mu_A = 0.626$. Using a higher value of β_{r_R} relative to β_{L_R} estimated by instrumenting for TransMilenio's location (Table 6, column 3), floorspace prices rise by 4.78% in the closed city model and 15.47% in the open city model. However as row 3 of Table A.1 shows, the difference does not matter for welfare. Inspection of the system of equations determining $\hat{\bar{U}}$ in the proof of Proposition 1 shows this depends only on the change in residential populations and employment.

Robustness. Table 9 explores the robustness of the welfare results to alternative parameter values. Panel A shows the results are highly robust to alternative estimates of the CMA elasticities from the IV specification, as well an alternative way to aggregate mode-specific commute times. The estimate of θ does matter quantitatively–the welfare effect falls from 2.21% to 1.15% as θ rises from the baseline estimate 3.398 to 6.15 estimated using IV–but the benefits from TransMilenio remain qualitatively large.[47] I retain the value of θ estimated via PPML as the baseline value given it accounts for zero commute flows in the data and is in line with existing estimates. Panel B shows that the welfare impact of TransMilenio is not very sensitive to alternative values of calibrated parameters.

Incorporating Congestion. While speeds for cars and other buses did not change on routes adjacent to TransMilenio (see Appendix F), the BRT could have had aggregate effects on road speeds that do not appear in a difference-in-difference specification. Appendix D.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 govern 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 both on new physical infrastructure as well as 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 Table 10. The left hand panel uses the average 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.[48] The main result is that allowing for congestion leads to a 30% larger welfare gain. This occurs due to the fact that as commuter substitute away from cars and buses and onto the BRT, they decongest the roads and reduce travel times on these other modes. The gap increases to 44% when using a larger value for the congestion elasticity of 0.097 from Allen and Arkolakis (2021) (estimated for intercity road travel in the US). 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.

5.3 Policy Counterfactuals

Impact of Alternative Networks. The first panel of Table 11 analyzes the impact of alternative network configurations. The first two rows simulate the effect of removing lines A and H that connect

---

[47] The IV estimate is reported in Table A.4. Ideally I would run the PPML model instrumenting for TransMilenio's location instead of the linear IV model (which drops observations equal to zero). However the IV-PPML model did not converge in the two-period specification, and has not been shown to yield consistent estimates. In general, non-linear models suffer from the incidental parameters problem (with the poisson model as an exception see Fernandez-Val and Weidner 2016).

[48] The welfare effect in this model with the congestion elasticity set to zero differs slightly from the baseline model since (i) the congestion elasticity is used when calibrating the unobserved traffic matrix for the observed equilibrium, and (ii) the construction of commute times is slightly different due to the routing model of commutes.

the North and South of the city with the CBD. The welfare impact of each line compared to the full network is 40% and 33% for the lines to the North and South respectively, with slightly larger impacts on output.

However, the welfare gains for these trunk lines are exceeded by the benefits from the feeder bus network (49% of the gains from the full network).[49] 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 a feeder system 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.

Land Value Capture One main criticism of TransMilenio was that it was not accompanied by an adjustment of zoning laws to allow housing supply to respond where it was needed. Appendix F 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 (see Hong et. al. 2015 for a review). These policies achieve the dual aim of increasing housing supply and raising revenue to finance the infrastructure's construction.

I evaluate the impact of TransMilenio if housing supply had responded to the opening of the system. As a benchmark, I allow housing supply to adjust to the increasing in floorspace values according a log-linear supply curve. Given 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 Oakland, CA, the 6th most inelastic city in the US from 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 Asian, European and American cities.[50] 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 D.4. I compare the two equilibria from first removing TransMilenio (without housing adjustment) and then adding it back under each housing supply model.

The last two panels of Table 11 presents the results. Panel B shows the impacts on welfare. Under free adjustment, welfare would have been 22.11% higher than it is today. Under the LVC schemes

---

[49] These gains do not have to sum to 100%, since these are statements about the impacts of removing components of the network leaving the others unchanged. Clearly, the value of a feeder system is conditional on the rest of the network being built.

[50] Compared to raising property taxes, these schemes are less likely to incur less opposition from stakeholders, are less distortionary, are more likely to work in settings with weak property tax systems, and provide additional benefits such as new residential and commercial units. See Hong et. al. (2015) and Salon (2014) for further details. My choice of parameters for this policy is 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, not dissimilar to my results.

welfare would have been 4.84% and 18.73% higher than it is today under the distance- and CMA-based policies respectively (with similar relative effects on city output). These welfare improvements come from increasing housing supply where it is demanded most as a result of new infrastructure, 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 most needed deliver bigger benefits. Panel C shows the fiscal benefits of LVC schemes. The distance-based instrument recoups between 9-14% of construction costs, while the CMA-based scheme covers 25-39% of costs (depending on how much the city population grows in response to the BRT).

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 adjustment 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 impact of new transit infrastructure in cities based on the change in accessibility induced by the infrastructure as well as the elasticities of economic activity to accessibility. Second, it shows these accessibility terms can be measured from readily available data, and estimates the elasticities 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 with the VTTS to isolate the importance of reallocation and general equilibrium effects in valuing the gains from new transit infrastructure.

The paper finds that the quantitative urban model performs well in explaining the heterogeneous adjustment of economic activity in the city to transit infrastructure, with the log-linear relationships between changes in economic outcomes and accessibility being borne out in the data. The VTTS only account for around 57.5% of the total welfare gain from the new transit infrastructure. Accounting for equilibrium effects therefore matters for valuing the gains from new transit infrastructure in cities, and the framework developed in this paper provides a blueprint to do so.

The paper 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 that passengers into the system's terminals have high returns. The second is that the welfare gains would have been around 19% higher had the the government implemented a more accommodative zoning policy, and government revenues from a land value capture 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.

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.

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 Sciency Policy & Practice*, 4(4): 495-508.

Baltes, M., J. Barrios, A. Cain, G. Darido, P. Rodriguez (2006), "Applicability of Bogotá's TransMilenio 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.

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.

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/t/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," Cadastre of Bogota 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," Cadastre of Bogota 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 Model Estimates

	(1)
Time	-0.013***
	(0.003)
Bus	-0.087**
	(0.041)
Car	-1.531***
	(0.191)
TM	-0.233***
	(0.072)
λ	0.157***
	(0.040)
Time of Day Controls	X
Demographic Controls	X

Note: Table shows estimation from nested logit regression on mode choices from trip-level data from the 2015 Mobility Survey. λ is the correlation parameter for the public nest. Bus, Car and TM show the mean preference term $b_m \equiv -\tilde{b}_m/\lambda_{k(m)}$ for each mode relative to walking. Demographic controls include a sex dummy as well as dummies for quintiles of the age distribution, time of day controls include dummies for the hour of trip departure. Each have choice-varying coefficients. Only trips to work during rush hour (hour of departure 4-8am) by individuals 18-55 are included. Robust standard errors are reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 2: Gravity Equation Estimates

	PPML	PPML	OLS	OLS	OLS, Wght	OLS, Wght
ln Commute Cost	-0.036**	-0.039**	-0.028	-0.035*	-0.042**	-0.046***
	(0.017)	(0.016)	(0.020)	(0.020)	(0.019)	(0.018)
N	710	710	576	576	576	576
Origin X Destination FE	X	X	X	X	X	X
Origin X Year FE	X	X	X	X	X	X
Destination X Year FE	X	X	X	X	X	X
Controls X Year FE		X		X		X

Note: Outcome is the commute shares in levels (PPML) or logs (OLS). Observation is an origin-destination-year cell. Only trips to work during rush hour (hour of departure 4-8am) by individuals 18-55 are included. Data is from 1995 and 2015 mobility surveys. Columns 1-2 estimate PPML models, 3-4 OLS models, and 5-6 OLS models weighted by the number of observations in each cell. Route-level controls are (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 3: Baseline Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Residents							
ln(Res Floorspace Price)	0.370**	0.382***	0.384***	0.342**	0.379***	0.228***	0.365**
	(0.147)	(0.143)	(0.141)	(0.146)	(0.123)	(0.075)	(0.171)
N	2,201	2,201	2,201	2,166	2,199	2,201	2,201
R^2	0.41	0.43	0.43	0.43	0.43	0.43	0.43
ln(Res Population)	0.722**	0.648**	0.746**	0.630**	0.683**	0.318**	1.086***
	(0.307)	(0.308)	(0.304)	(0.302)	(0.269)	(0.160)	(0.349)
N	2,256	2,256	2,256	2,219	2,255	2,254	2,256
R^2	0.34	0.35	0.37	0.37	0.37	0.37	0.37
Panel B: Firms							
ln(Comm Floorspace Price)	0.526**	0.540**	0.621***	0.580**	0.514**	0.359***	0.718**
	(0.232)	(0.232)	(0.230)	(0.232)	(0.199)	(0.125)	(0.282)
N	2,080	2,080	2,080	2,047	2,089	2,084	2,080
R^2	0.09	0.10	0.11	0.11	0.11	0.11	0.11
Comm Floorspace Share	0.290***	0.297***	0.291***	0.284***	0.197***	0.152***	0.286***
	(0.079)	(0.081)	(0.080)	(0.081)	(0.066)	(0.043)	(0.094)
N	2,230	2,230	2,230	2,195	2,239	2,234	2,230
R^2	0.14	0.15	0.15	0.15	0.15	0.15	0.15
ln(Establishments)	2.101***	1.787***	1.168*	1.114*	1.266**	0.847***	1.238
	(0.611)	(0.619)	(0.604)	(0.608)	(0.517)	(0.322)	(0.764)
N	2,028	2,028	2,028	1,995	2,028	2,028	2,028
R^2	0.21	0.23	0.27	0.27	0.27	0.27	0.27
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 establisment 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. Standard errors clustered by census tract reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 4: Planned Networks

	(1)	(2)
Panel A: Residents		
ln(Res Floorspace Price)		
Δ ln(CMA)	0.384***	0.368**
	(0.141)	(0.147)
E[Δ ln(CMA Plan)]		0.084
		(0.222)
N	2,201	2,201
R^2	0.43	0.43
p-val		0.70
ln(Res Population)		
Δ ln(CMA)	0.746**	0.817***
	(0.304)	(0.311)
E[Δ ln(CMA Plan)]		-0.388
		(0.355)
N	2,256	2,256
R^2	0.37	0.37
p-val		0.27
Panel B: Firms		
ln(Comm Floorspace Price)		
Δ ln(CMA)	0.621***	0.687***
	(0.230)	(0.243)
E[Δ ln(CMA Plan)]		-0.360
		(0.431)
N	2,080	2,080
R^2	0.11	0.11
p-val		0.39
Comm Floorspace Share		
Δ ln(CMA)	0.291***	0.290***
	(0.080)	(0.081)
E[Δ ln(CMA Plan)]		0.005
		(0.099)
N	2,230	2,230
R^2	0.15	0.15
p-val		0.96
ln(Establishments)		
Δ ln(CMA)	1.168*	1.005*
	(0.604)	(0.606)
E[Δ ln(CMA Plan)]		0.847
		(0.636)
N	2,028	2,028
R^2	0.27	0.27
p-val		0.18

Note: Column (1) repeats the baseline specification i.e. column (3) from Table 3. 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 Δ ln(CMA) in columns 1 and 2. Standard errors clustered by census tract reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: Staggered Station Openings

	(1)	(2)	(3)	(4)
Panel A: Residents				
	Res Pr	Res Pr	Res Pop	Res Pop
ln(RCMA)	0.183**	0.231**	0.396*	0.521**
	(0.088)	(0.097)	(0.216)	(0.239)
ln(RCMA) Later Phase		0.336		0.345
		(0.216)		(0.317)
N	2,144	2,144	2,207	2,207
R^2	0.45	0.45	0.29	0.29
Panel B: Firms (Land Markets)				
	Comm Pr	Comm Pr	Comm Sh	Comm Sh
ln(FCMA)	0.483***	0.478**	0.210***	0.206***
	(0.186)	(0.187)	(0.054)	(0.054)
ln(FCMA) Later Phase		0.281		0.220
		(0.614)		(0.175)
N	2,055	2,055	2,182	2,182
R^2	0.06	0.06	0.08	0.08
Panel C: Firms (Census Employment)				
	Emp	Emp	Form Emp	Form Emp
ln(FCMA)	1.662	1.835*	2.217*	2.285*
	(1.051)	(1.057)	(1.341)	(1.351)
ln(FCMA) Later Phase		1.717		0.647
		(1.327)		(1.779)
N	1,927	1,927	1,629	1,629
R^2	0.22	0.22	0.16	0.16

Note: Table repeats the baseline specification i.e. column (3) from Table 3. 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 in 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. Standard errors clustered by census tract reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 6: IV Estimates

	Baseline	IV	IV-LCP&Tram
Panel A: Residents			
ln(Res Floorspace Price)	0.384***	0.276***	0.951***
	(0.141)	(0.102)	(0.223)
N	2,201	2,201	2,172
F-Stat		2,475.45	283.08
ln(Res Population)	0.746**	0.553**	0.637
	(0.304)	(0.225)	(0.453)
N	2,256	2,256	2,214
F-Stat		2,404.19	412.04
Panel B: Firms			
ln(Comm Floorspace Price)	0.621***	0.552***	0.560*
	(0.230)	(0.204)	(0.304)
N	2,080	2,080	2,053
F-Stat		3,165.71	743.36
Comm Floorspace Share	0.206**	0.257***	0.213**
	(0.084)	(0.071)	(0.105)
N	2,231	2,230	2,200
F-Stat		3,112.12	715.06
ln(Establishments)	1.168*	1.082*	1.251
	(0.604)	(0.560)	(0.786)
N	2,028	2,028	1,965
F-Stat		2,854.69	487.47

Note: Observation is a census tract. Specification corresponds to column (3) of Table 3. 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 intruemnts. 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). Standard errors clustered by census tract reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 7: Welfare Results: First Order vs General Equilibrium

First Order Approximation (VTTS)	
Welfare Gain (%)	1.302
General Equilibrium (Main Estimates)	
Welfare Gain (%)	2.262
90% CI	(0.701,5.214)
95% CI	[0.527,6.748]
Fraction of Gains from VTTS	57.573
General Equilibrium (No Externalities)	
Welfare Gain (%)	1.527
Fraction of Gains from VTTS	85.289

Note: Table shows the percentage change in welfare from 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. The first panel is the first order welfare approximation using the VTTS from Proposition 2. The change in travel times accounts for the discrete choice over modes used to aggregate mode-specific travel times as in Section 4.2. The second panel is the full general equilibrium response using the main estimates in column (3) of Table 3 of the CMA elasticities. 90% and 95% confidence intervals are provided by bootstrapping the quantitative exercise 200 times as described in Appendix Section C.8. The last entry shows the fraction of gains accounted for by VTTS. The third panel is the full general equilibrium response in a model without externalities. I compute the reduced form elasticities using the expressions derived in Appendix Section C.1, using my 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 and thus is not comparable with the second panel. The last entry again shows the fraction of gains accounted for by VTTS in this model.

Table 8: Main Quantitative Results

	No Migration	Migration
Panel A: Aggregate Effects (%)		
GDP	2.98	12.71
Welfare	2.21	0.50
Population	0.00	7.97
Rents	-0.71	4.33
Panel B: Observed Growth Accounted For (%)		
GDP Growth	2.83	12.06
Population Growth	0.00	29.24
Panel C: Costs & Net Benefits (mm 2016 USD)		
NPV Increase GDP	41,620.96	177,599.26
Capital Costs	1,449.75	1,449.75
NPV Operating Costs	7,180.53	7,180.53
NPV Total Costs	8,630.28	8,630.28
NPV Net Increase GDP	32,990.67	168,968.98
% Net Increase GDP	2.36	12.09

Note: Panel A 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 Section D.1 for details). Panel B shows 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. Panel C shows the costs and net benefits, computing net present values (NPV) over a 50 year time horizon with a 5% interest rate. All numbers are in millions of 2016 USD. The NPV of the increase in GDP is simply the NPV of the change in Bogotá's GDP in dollar values. Capital costs are the one-time infrastructure costs of building the network. Total costs are the one-time capital costs associated with building the network combined with the NPV of operating costs. The NPV net increase in GDP nets this out from the gross gains in the first row, while the final row converts this back into a fraction of 2016 GDP. See Appendix Section E.4 for more details. Note the average welfare value differs from Table 7, which uses the counterfactual equilibrium without TransMilenio as the initial equilibrium for ease of comparison with the VTTS.

Table 9: Change in Welfare (%): Robustness

	No Migration	Migration
Panel A: Alt. Estimated Params		
Baseline	2.21	0.50
IV	2.17	0.34
IV-Loc	2.25	0.46
Alternative Times	2.00	0.16
θ OLS	1.87	0.43
θ IV	1.15	0.27
Panel B: Alt. Calibrated Params		
$\sigma = 4$	2.45	0.87
$\sigma = 8$	2.11	0.41
$\beta = 0.8$	2.44	0.58
$\beta = 0.7$	2.13	0.48
$\rho = 6$	2.21	0.86

Note: Table shows the percentage change in average welfare (as defined in Table 8) under alternative parameter values. Panel A examines sensitivity to alternative values of estimated parameters. The first row recreates the baseline results. The second row uses the CMA elasticities from the second column of Table 6 which instrument for the realized change in CMA (i.e. the term that does not hold residential population and employment fixed at their initial value in the post-period) using the baseline measure. The third row uses the CMA elasticities from the third column of Table 6 when instrumenting for the realized change in CMA using the LCP and Tram instrument. The fourth row uses the coefficients from column 6 of Table 3, using an alternative method to aggregate mode-specific commute times. The fifth row uses an alternative value for $\theta = 3.97$ estimated via OLS in column 4 of Table 2. The six row uses a value for $\theta = 6.15$ estimated via IV using the LCP and Tram instrument in column 3 of Table A.4. Panel B varies the value of calibrated parameters.

Table 10: Incorporating Congestion

	Baseline Congestion Elasticity		**Alternative Congestion Elasticity**	
	% Change in Welfare	% of Welfare Change without Congestion	% Change in Welfare	% of Welfare Change without Congestion
No Congestion	2.410	100.00	2.367	100.00
Congestion	3.117	129.37	3.405	143.83
Convert TM to Car Lanes	0.014	0.59	0.023	0.95

Note: Table reports welfare results from model allowing for congestion (see Appendix D.2 for details). The first panel uses the baseline congestion elasticity of 0.06, the average congestion elasticity estimated for Bogotá by Duranton and Akbar (2017). The second panel uses an alternative value congestion elasticity of 0.0972 from Allen and Arkolakis (2021). 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 of 2.21% 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. This exceeds the baseline impact since the substitution of individuals from buses and cars reduces congestion and travel times on roads. The third row shows the welfare impact had TransMilenio routes been made into car lanes instead of BRT (the absolute value of $\bar{U}^{NoTM}/\bar{U}^{ReplaceTMWithRoads} - 1$). Odd columns show the welfare effect, while even columns show the welfare effect as a percent of the impact in the model without congestion in row 1. The "No Congestion" change in welfare under the alternative congestion elasticity differs from the value under the baseline elasticity since the elasticity is used when calibrating the (unobserved) traffic matrix in the initial equilibrium.

Table 11: Policy Counterfactuals

Panel A: Alternative Networks

	Fraction of Gains from Full Network	
	Welfare	**Output**
Line South	40.03	37.23
Line North	33.48	34.40
Feeders	48.53	37.23

Panel B: Land Value Capture Welfare Effects

	% Increase Relative to Baseline	
	Welfare	**Output**
Free Adjustment	22.11	8.21
LVC, Bands	4.84	2.85
LVC, CMA	18.73	6.28

Panel C: Land Value Capture Revenue Effects

	Closed City	**Open City**
LVC Band Revenue (mm)	132.484	212.093
As share of capital costs	9.14	14.63
LVC CMA Revenue (mm)	375.798	572.291
As share of capital costs	25.92	39.48

Note: Panel A shows the impact of particular network components relative to the full network using the baseline model. Each row reports the percentage change in each variable from a simulation that removes each system component, as a fraction of the percentage change in each variable under removal of the full network. Panel B shows the impacts of alternative housing supply models, using the model extension from Section D.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

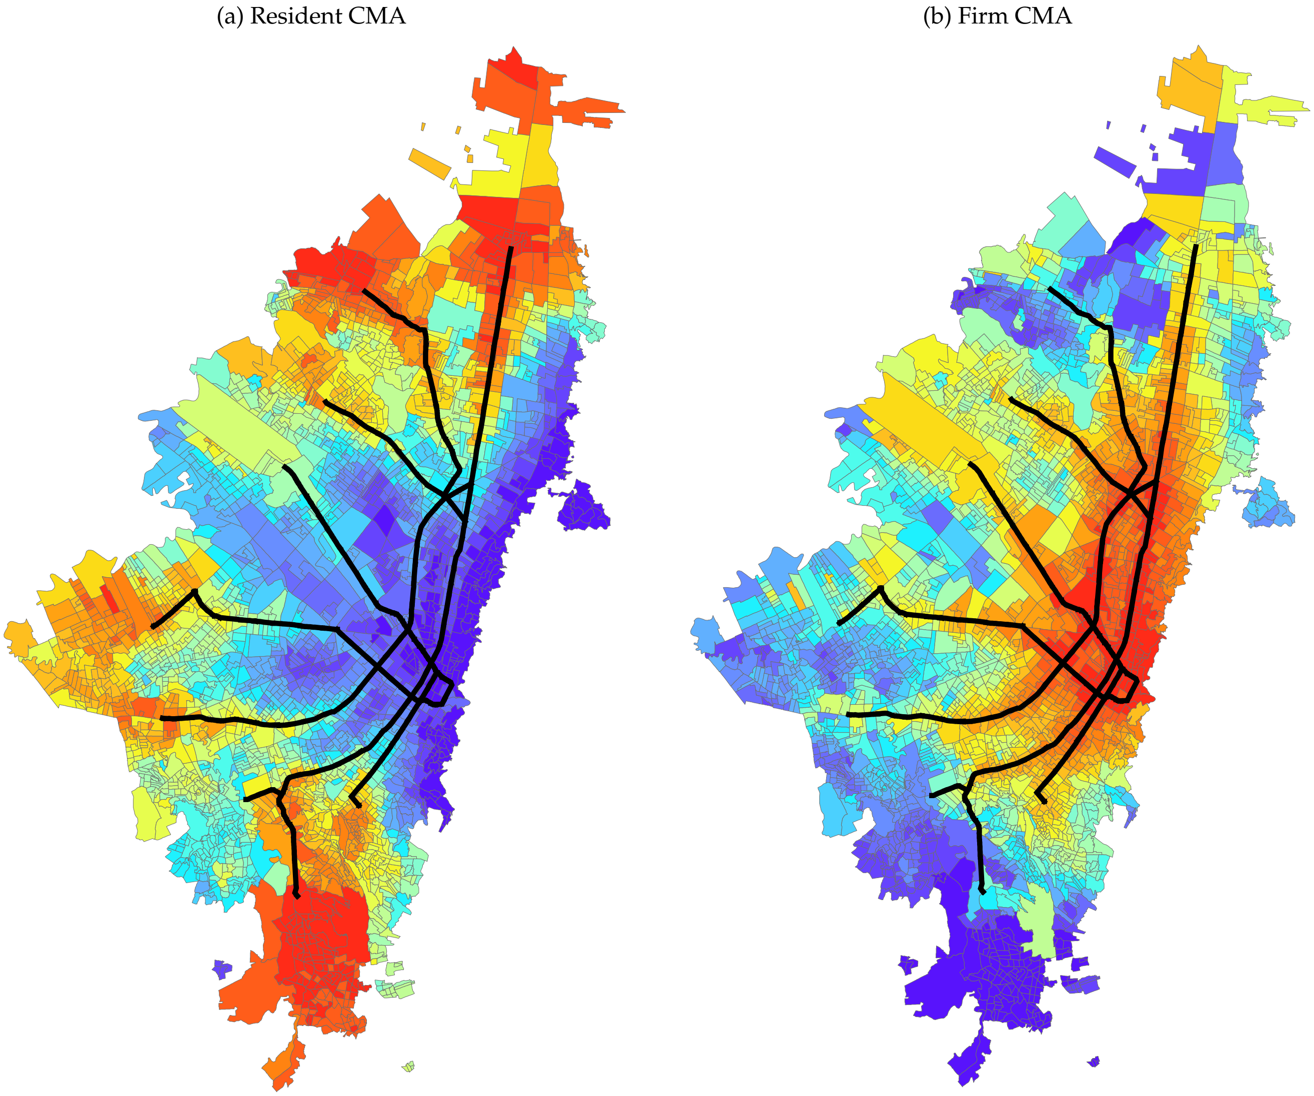

(a) Resident CMA

(b) Firm CMA

Note: Plot shows the change in CMA from the baseline specification. Population and employment arefixed 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 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

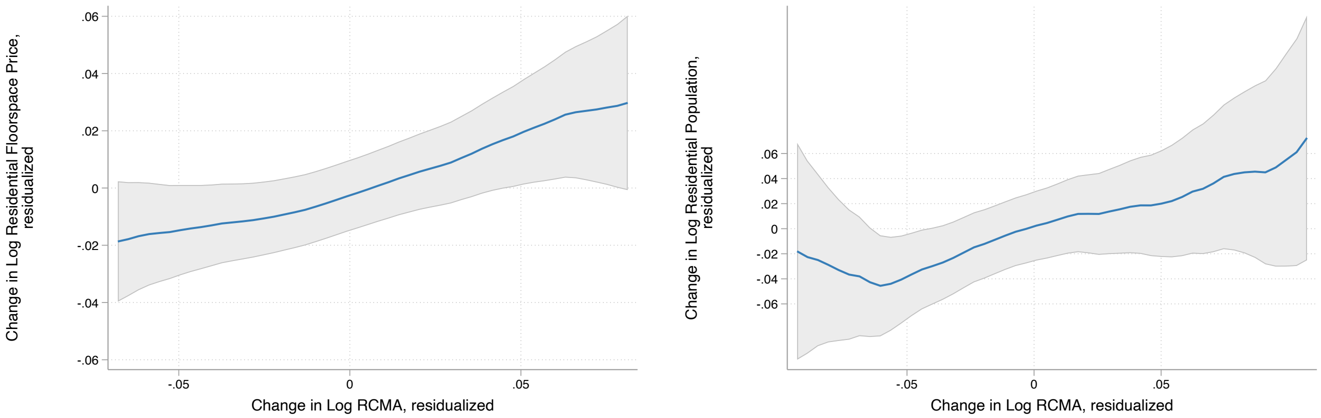

(a) Residential Floorspace Prices

(b) Residential Population

(c) Commercial Floorspace Prices

(d) Employment

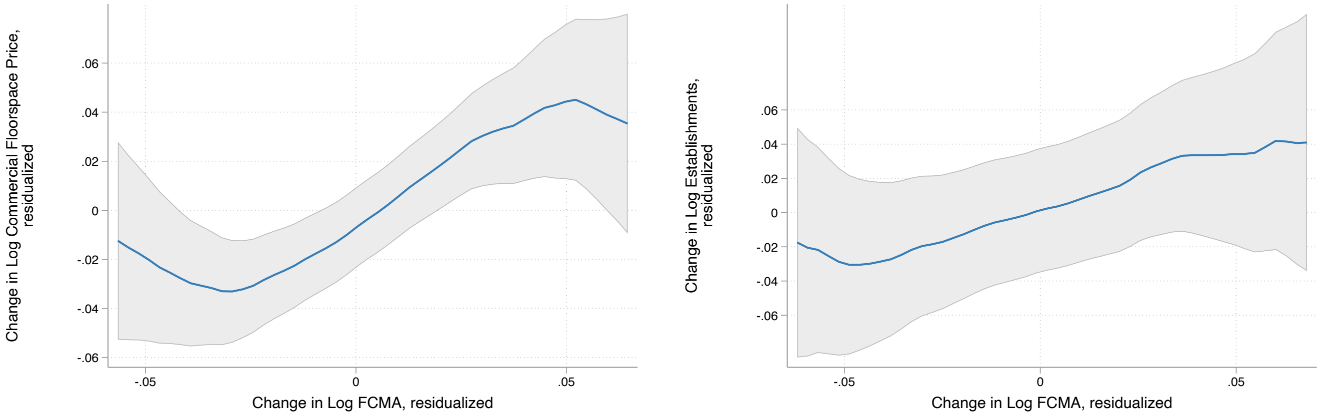

Note: Plot shows the non-parametric relationship between outcomes and CMA. Specifications correspond to the reduced form from column (3) of Table 3. Top and bottom 2% of the change in CMA are trimmed to reduce noise at the tails and zoom in on main relationship.

Figure 3: Residential Floorspace Price Event Study

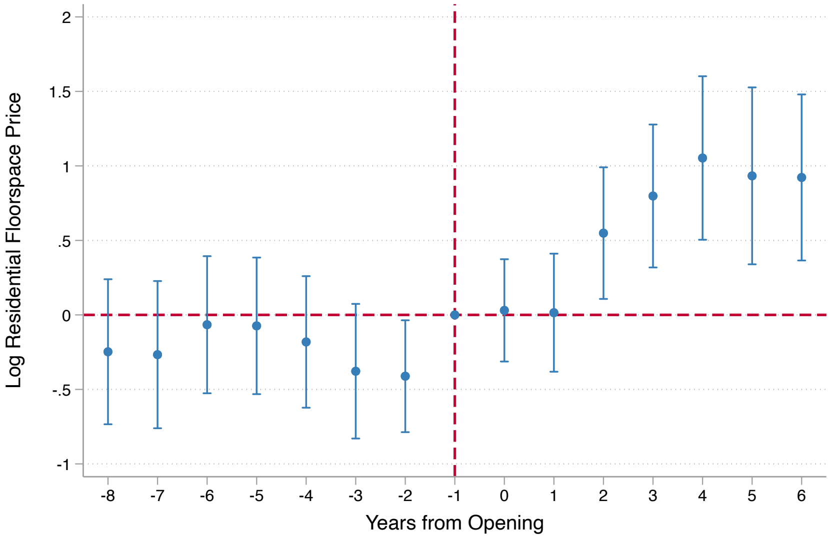

Note: See discussion in Section 4.3 for details. The year before opening is the omitted category.