

The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio*

Nick Tsivanidis[†]
Dartmouth College

February, 2019

Abstract

How does public transit infrastructure shape the structure of cities, and how are the gains shared between low- and high-skilled workers? The standard approach in transportation economics measures the benefits in terms of the value of time saved, but new infrastructure can change individuals' decisions of where to live and drive land and labor market adjustment. This paper develops a quantitative general equilibrium model of a city where low- and high-skill workers with non-homothetic preferences sort over where to live, where to work, and whether or not to own a car. This theory provides a new reduced form framework to evaluate the effects of transit based on "commuter market access". I leverage detailed tract-level data spanning the construction of the world's largest Bus Rapid Transit system—TransMilenio—in Bogotá, Colombia, to show this explains the change in city structure in response to the new infrastructure. While the system caused increases in welfare and output larger than its cost, the high-skilled benefit slightly more.

*Thanks to Chang-Tai Hsieh, Marianne Bertrand, Erik Hurst and Robert E. Lucas Jr. for their guidance and support. I also thank Rodrigo Adao, Fernando Alvarez, Jonathan Dingel, Rick Hornbeck, Ralph Ossa, Felix Tintelnot, Nancy Stokey, Owen Zidar, my discussants Treb Allen, David Atkin and Cecile Gaubert, and participants of working groups at Berkeley (ARE/Econ/Haas), Boston College, Chicago, Chicago Fed, CMU Tepper, Dartmouth, Harvard/MIT, IGC Cities, Johns Hopkins SAIS, LBS, LSE, Minneapolis Fed, NBER SI (ITI/DEV), NYU, Stanford (Econ/GSB), St Louis Fed, UMD College Park, Universidad de los Andes, Universidad de Javeriana, Urban Economics Association, USC, Wharton, World Bank and Yale 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. 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.

[†]Department of Economics, Dartmouth College. Email: nick.tsivanidis@dartmouth.edu. First version: November 2017.

1 Introduction

How large are the economic gains to improving public transit systems within cities and how are they distributed between low- and high-skilled workers? With 2.5 billion people predicted to move into cities by 2050, mostly in developing countries, governments will spend vast sums on mass transit systems to reduce congestion associated with this rapid urban growth.¹ The reliance of poor, low-skilled individuals on public transit suggests they may benefit the most. Yet measuring the benefits of these systems is challenging: while individuals save time on any particular commute, their decisions of where to live and work will change as new alternatives become attractive and land and labor markets adjust. 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 before and after 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 aggregate and distributional effects of urban transit systems. First, I build a quantitative general equilibrium model of a city where low- and high-skill workers with non-homothetic preferences sort over where to live, where to work, and whether or not to own a car. Second, I show that in a wide class of models the effect of the entire transit network on a location is summarized by a single measure: “commuter market access” (CMA). For individuals this reflects access to jobs while for firms it reflects access to workers; both are easily computed using data on employment and residence across the city. A special case of my model admits a log-linear reduced form relationship linking outcomes such as population, employment and house prices to CMA, providing a regression framework to guide empirical work. Third, I use the change in commuting infrastructure provided by TransMilenio to assess the model’s ability to explain the change in city structure and quantify the effects of the system and counterfactual policies.

I have three main results. First, changes in CMA parsimoniously explain the heterogeneous adjustment of population, employment and housing markets to TransMilenio. This suggests the framework can be applied elsewhere to improve predictions about the effects of transit on the spatial organization of cities. Second, I find the system provided large aggregate gains for the city, increasing average welfare by 3.5% and output by 2.73% (net of construction and operating costs) at my most conservative estimates. However, these gains would have been around one fourth larger had the government implemented a complementary change in zoning policy to allow housing supply to respond where it was most needed. Third, I find that high-skilled workers benefitted slightly *more*, which is surprising given the reliance of the low-skilled on public transit.

To build intuition, I find certain key channels explain the incidence of improved public transit across worker groups. The first is mode choice: the group that relies on public transit benefits more. This operates in favor of the low-skilled who are poorer and less willing to pay for cars. The second is the elasticity of commuting decisions to commute costs, which determines how willing individuals are to bear high commute costs to work in a particular destination. In the model, this is determined by the heterogeneity of workers’ match-

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

productivities with firms in different locations. For example, a high-skilled IT worker may be more willing to incur a costly commute to an especially well-paid position. A low-skilled cleaner who receives similar wages wherever they work may instead substitute towards other alternatives. In the data, I find high-skilled workers' commuting choices are less sensitive to relative differences in commute costs - suggesting a greater dispersion of match productivity, consistent with other evidence - so they bear a greater incidence along this channel.² Lastly, the geography of the city and transit network matter: where house prices appreciate and wages adjust, whether the system connects locations of dense residence with well-paid jobs, and how these characteristics differ where each group lives and works. In Bogotá these favor the high-skilled.

Since this last channel is context-specific, to draw a more generalizable result I compare my results with those using the standard approach in transportation economics. This evaluates the gains from new infrastructure based on the value of travel time saved, often used by institutions such as the World Bank. I show this is precisely the first order welfare effect in an efficient equilibrium of the quantitative urban model. While this approximation only slightly overstates the level of equilibrium welfare gains, it delivers the opposite result for the distributional impacts with a large decrease in inequality. This is because it focuses on mode choice alone, without accounting for the spatial reorganization of the city and general equilibrium adjustment of prices that are important for large shocks. My findings therefore imply that investments in public transit are a less precise way to target welfare improvements for the poor than is implied by the travel time savings approach.

Opened in 2000, TransMilenio is the world's most used BRT system with a daily volume of over 2.2mn 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, passengers pay at station entrances using smart cards, and buses are boarded at stations rather than at roadside. BRT provides an attractive alternative to subways in rapidly growing developing country cities since they are able to 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 2,800 census tracts on residence, employment, commuting patterns, and land markets spanning the system's construction.

Prior to TransMilenio, low-skilled workers commuted using a network of informal buses which were on average 30% slower than cars. To understand the implications of improving public transit on worker welfare, I develop a quantitative general equilibrium model of a city where workers choose where to live, where to work, and how to commute. Non-homothetic preferences mean that the high-skilled live in high amenity neighborhoods and are more likely to own cars. Individuals work in different locations due in part to differential demand for skills from firms across the city. For example, retail and manufacturing establishments demand more low-skilled workers while real estate and financial service businesses rely on the high-skilled. Individuals differ in their match-productivity with firms in each location and their preference to live in each neighborhood. Together these determine the sensitivity of commute flows to commute costs. Differences in residential loca-

²See Lee (2015) for an estimates of this dispersion across education groups using wage data in the US and developing countries. One might expect rich, high-skilled workers to be more sensitive to commute costs since their value of time (VoT) is higher. Indeed, in the model VoT is proportional to wages and it is precisely because of this that the high-skilled are more willing to pay for faster transit (i.e. cars). However, when individuals decide where to work, they choose based on *relative* differences in wages net of commute costs across locations. Thus, the sensitivity of commute choices to commute costs depends on the dispersion rather than level of wages.

³For example, the per mile cost of construction of the subway in Colombia's second largest city, Medellín, was ten times that of TransMilenio, all the while maintaining similar system speeds. Moreover, TransMilenio took less than eighteen months to construct and open, compared to the twelve years taken by Metro Medellín.

tions, commuting elasticities and the relative demand for worker skills turn out to be crucial in determining the distributional effects from improving transit.

A large literature estimates average treatment effects of transit based on proximity to stations. In contrast, I show that for a wide class of models featuring a gravity equation for commute flows the full direct and indirect effects of the entire transit network on firms and workers can be summarized by a single variable: CMA. Importantly, these terms are easily computed using data on residence and employment in the city (available from censuses, or alternative sources such as cellphone metadata records), as well as a measure of commute costs. Figure 1 plots the change in CMA as a result of TransMilenio. For residents, this captures access to high paying jobs through the commuting network. Tracts towards the edge of the city far from the high density of jobs in the center experienced a much larger improvement in market access. For firms, it reflects access to workers. Central locations benefit most from increased access to workers supplied along all spokes of the network. Changes in market access capture a wide heterogeneity in treatment effects from TransMilenio (separately for workers and firms) that would be missed by looking at distance to the system alone.

In a special case of my model, the equilibrium has a reduced form in which outcomes such as population, employment and house prices can be written as log-linear functions of CMA. Moreover, I show a class of models with log-linear demand for residents and workers across the city have a similar representation. The framework is therefore isomorphic to a number of alternative assumptions over production technologies, housing supply and worker preferences. I use this regression framework as a model validation exercise to examine whether the market access approach explains the change in city structure resulting from the new transit infrastructure. I also use the estimated elasticities to approximate the first order welfare gains from the system.

To address non-random route placement, I predict TransMilenio's location in two ways. First, I use a historical tram system built by 1921. Second, using engineering estimates for the cost of building BRT on different types of land use, I solve for least-cost construction paths connecting terminals at the end of the system with the central business district (CBD) as was the intent of the government. I then construct instruments for the change in CMA had TransMilenio been built along these predicted routes.

My identification assumption is that these instruments have only an indirect effect on outcomes through the probability of TransMilenio being built. 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. Relative to distance-based analyses, a key advantage of my approach is that I can control for the distance to these features and use only residual variation in predicted CMA growth for identification. To provide additional evidence the effects are causal I run falsification tests exploiting the timing of station openings, use residual variation in market access conditional on distance to stations and leverage changes in CMA to locations further than 1.5km from each tract.

I find that changes in CMA perform well in explaining the heterogeneous response of population, employment and land markets in response to TransMilenio. Improvements in resident CMA led to growth in commute distances and wages, supporting the intuition that it measures access to jobs. The system also caused a re-sorting of workers by skill group. The high-skilled moved into high-amenity, expensive neighborhoods in the North while the low-skilled moved into poorer neighborhoods in the South. This suggests that transit has the

potential to increase residential segregation between skill groups in cities.⁴

In the final part of the paper, I structurally estimate the full (non-linear) model. Some parameters, such as spillovers in productivities and amenities, are challenging to estimate in cross-sectional data. For example, a location's productivity may be a cause or consequence of the number of workers employed there. Since the supply of workers and residents in the model is a log-linear function of market access, my instruments provide exogenous variation in the number of individuals living and working across the city. This allows me to identify these key elasticities through a Generalized Method of Moments (GMM) procedure.

I estimate an agglomeration elasticity roughly three times the size of median estimates in the US but close to other studies using experimental approaches. I provide one of the first estimates identified using variation within a less developed country city, suggesting that these forces can be particularly strong in poorer countries. I find a substantial elasticity of amenities to the college share of residents, reflecting the endogeneity of neighborhood characteristics like crime.

The model performs well in matching a number of non-targeted moments such as income, employment and commute flows by skill group, and the change in residential segregation. The amenities and productivities recovered from the model correlate well with observable proxies like local homicide rates and the slope of land. I check the robustness of my results to alternative parameter values and incorporate home ownership, alternative timing assumptions and the employment of domestic servants in model extensions. Lastly, I use the extreme assumptions of either zero or infinite mobility costs between Bogotá and the rest of Colombia to bound the impact of TransMilenio on welfare, population, land rents and output.

The system led to large aggregate gains in worker welfare and output. Productive locations were able to “import” more workers through the commuting network. This suggests better transit can improve the spatial allocation of labor within cities. The increase in output greatly exceeded construction and operating costs, supporting the notion that BRT can be a profitable investment. Population decentralized, as improved access to jobs made distant neighborhoods more attractive. Land use became more specialized: the share of floorspace used for commercial (residential) purposes increased in central (outlying) locations where firm (resident) commuter access increased the most.

High-skill workers benefit slightly more from the system. The incidence of public transit across skill-groups is determined not only by who uses it most, but also by how easily individuals substitute between commutes, whether the system connects workers with employment opportunities, and equilibrium adjustment of housing and labor markets. Landlords benefit from house price appreciation where transit access improves.

The effect of different parts of the network is heterogeneous: lines serving poor neighborhoods in the South of the city, as well as a cable car connecting hillside slums with a TransMilenio, disproportionately benefit the low-skilled. The conclusion that the low-skilled benefit less than what is implied by mode choice alone remains generalizable, though, since existing evidence suggests that the key elasticities that vary across groups have similar relative magnitudes in other countries. Moreover, the methodology developed in this paper can be applied to the specific geography and transit systems of any city where data on residence and employment are available to evaluate the effects of new infrastructure.

⁴Heilmann (2016) notes a similar finding following the opening of the light rail system in Dallas.

I use the model to assess the impact of counterfactual policies. First, I simulate the removal of the feeder bus network that transports individuals in the outskirts of the city to terminals at the end of lines using existing road infrastructure at no additional fare. This part of TransMilenio increases welfare more than any other single line of the network. This underlines the potential for large benefits to providing cheap, complementary services that reduce the last-mile problem of traveling between stations and final destinations.

In the second exercise, I evaluate the welfare impacts of a “Land Value Capture” (LVC) scheme under which development rights to increase building densities near stations are sold by the government to developers, and the extent to which the revenues could have financed the system’s construction. Similar schemes have seen great success in Asian cities such as Hong Kong and Tokyo. In contrast, 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 where it was needed. I compare the effects of two alternative policies. The first increases permitted densities within a certain distance of stations, while the second allocates the same number of permits based on predicted growth in CMA. The CMA-based policy increases welfare gains from TransMilenio by around 23%, while government revenues cover at least 18% of the construction costs. Under the distance-based scheme welfare and government revenues only increase by around half that amount. These policies disproportionately benefit low-skilled workers by dampening house price appreciation towards the edge of the city where they live. My findings suggest large returns to the pursuit of an integrated transit and land use policy, and highlight the potential of CMA as an instrument to guide government policy.

The rest of the paper proceeds as follows. Section 2 discusses the paper’s contribution to the literature. Section 3 presents the context of Bogotá and TransMilenio. Section 4 develops the model and Section 5 outlines the reduced form framework it delivers. Section 6 describes the data. Section 7 presents the reduced form estimation results while Section 8 structurally estimates the full model which Section 9 uses to quantify the effects of TransMilenio. Section 10 simulates the effects of counterfactual policies. Section 11 concludes.

2 Relation to Previous Literature

This paper contributes to several literatures. Most closely related is the body of work that examines the impact of transportation infrastructure on economic activity. A first 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 2016; Baum-Snow et. al. 2017) or distance to stations (Gibbons and Machin 2005; Glaeser et. al. 2008; Billings 2011). However, when spatial units are interlinked spillovers across treatment and control locations confound causal inference from such comparisons.⁵ 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. I contribute to this literature by deriving a measure from a class of commuting models that explicitly captures the full direct and indirect effects of changes in the transit network between connected locations, allowing for a causal identi-

⁵Identification of causal impacts of transit infrastructure from relative comparisons requires that the treatment effect on one location is independent of whether another is treated. This is known as the stable unit treatment value assumption (SUTVA), and can be violated if locations interact which is likely in localized urban geographies. See Donaldson (2015) for a full discussion.

fication of transit connections that captures heterogeneous responses as a function of city geography.⁶ A second strand of this literature explores the effect of infrastructure between regions on economic development through access to goods markets (Redding and Sturm 2008; Donaldson forthcoming; Bartelme 2015; Donaldson and Hornbeck 2015; Alder 2019).⁷ However these models contain no notion of commuting within cities, and are silent on the effects of transit infrastructure. In contrast, I consider a different class of urban commuting models where individuals can live and work in separate locations, and show that reduced form relationships between outcomes and model-derived accessibility measures explain the change in city structure in response to a real world change in transit infrastructure. Lastly, my paper relates to work in transportation economics that measures the benefits of improved transportation through the value of travel time saved (Train and McFadden 1978; Small and Verhoef 2007), which I show is precisely the first order welfare effect in an efficient equilibrium of a quantitative urban model.

I also contribute to the growing body of work on quantitative spatial models (Ahlfeldt et. al. 2015; Allen et. al. 2015; Fajgelbaum and Schaal 2017; Monte et. al. 2017; Owens et. al. 2017; Severen 2017; Bryan and Morten 2018; Heblich et. al. 2018; Adao et. al. 2019; Allen and Arkolakis 2019). First, I incorporate multiple types of workers, firms and transit modes, which is necessary to assess the distributional impacts of improvements in particular modes of transit. My model includes non-homothetic demand for housing and car ownership, which I find to be important: a homothetic model is unable to match the sorting responses to improved accessibility observed in the data, underestimates the increase in inequality by about one third, and shrinks the estimates of amenity spillovers by about one quarter. Allowing for multiple firm types enables me to recover unobserved relative wages that vary by employment location, which provide a better fit to the distribution of employment across skill groups and have important equity implications. Second, I show these models share a common measure that summarizes the effect of transit on the supply of residents and workers across locations. These measures are easily computed using data and population and employment, and a class of models admit an exact reduced form representation in terms of these accessibility measures that can be used to evaluate the impact of changes in transit infrastructure. Third, I leverage the construction of the world's largest BRT in a validation exercise to show the regression framework provided by the model performs well in explaining the change in city structure.

Lastly, this paper connects with an extensive literature on agglomeration spillovers (see Rosenthal and Strange 2004 for a review). A smaller strand of work uses potentially exogenous sources of variation in the density of economic activity to estimate these spillovers (Combes et. al. 2010; Greenstone et. al. 2010; Kline and Moretti 2014; Ahlfeldt et. al. 2015). Other papers examine how amenities depend on the composition

⁶In the urban planning literature, the notion of "accessibility" as a determinant of land use and prices within cities dates back as far as Hansen (1959). The idea has received little attention within urban economics which instead has focused on models where distance from the CBD is a summary statistic for the spatial configuration of the city (Alonso 1964; Mills 1967; Muth 1969; Lucas and Rossi-Hansberg 2002). There is a literature that measures accessibility in polycentric cities (Handy and Niemeier 1997; Anas et. al. 1998) and that empirically relates it to residential property prices (Ahlfeldt 2011; McArthur et. al. 2012) in cross-sectional comparisons. In contrast, I show similar accessibility measures can be derived from general equilibrium theory separately for firms and workers, and use a change in the transit network to estimate their causal impact on a range of urban outcomes. My focus on accessibility links to recent work on urban form and consumption-related trips (Couture 2016; Duranton and Turner 2017).

⁷There is a long literature using an empirical market access approach in international trade literature e.g. Redding and Venables (2004). I also contribute to a growing literature on the impacts of BRT in developing country cities (Gaduh et. al. 2018; Majid et. al. 2018).

of local residents (Bayer, Ferreira and McMillan 2007; Guerrieri, Hartley and Hurst 2013; Diamond 2016; Giannone 2018). I contribute to this literature by providing intra-city estimates of productivity and amenity spillovers within a developing country city, identified using an expansion in the transit network that separately shifts the supply of labor and residents across the city.

3 Background

Bogotá is the political and economic center of Colombia, accounting for 16% and 25% of the country’s population and GDP respectively. Its population of eight million inhabitants makes it the world’s ninth densest, with a stark divide between rich and poor.⁸ In this section, I provide background on the city and its transit system.

3.1 Structure of Bogotá

Residence and Employment Bogotá is characterized by a high degree of residential segregation between the rich and poor. Defining high-skill workers as individuals who have completed some post-secondary education, panel (a) in Figure 2 plots the share of college residents within a census tract in 1993.⁹ The high-skilled are more likely to live in the North, with the low-skilled workers living in the city’s South and periphery. Panel (b) shows that these poorer neighborhoods have a much higher population density, reflecting the smaller per capita housing consumption.

High- and low-skilled residents work in different industries and neighborhoods. Table 1 shows the share of workers employed in each one-digit industry with post-secondary education. Workers in domestic services, hotels and restaurants, manufacturing and retail are relatively unskilled, while those in real estate, education and financial services tend to be high-skilled. These jobs are located in different parts of the city. Defining high-skill intensive industries as those with college employment shares above the median, Figure 3 shows that while overall employment is concentrated along two bands to the west and north of the city center, high-skill intensive industries are located more towards the North.

Commuting Prior to TransMilenio In 1995 the average trip to work in Bogotá took 55 minutes, more than double the average commute in US cities. The vast majority of these commutes was taken by bus (73%), followed by car (17%) and walking (9%).¹⁰ 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. However, 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

⁸Colombia is the eleventh most unequal country in the world according to the ranking of Gini coefficients from the World Bank. The income distribution in Bogotá has a slightly higher Gini than the country as a whole in 2014 (most recently available data).

⁹Datasets are described in Section 6. In this section, population data is from the 1993 census, employment location data uses the 1990 economic census, other employment data is from DANE’s GEIH and ECH and commuting data is from DANE’s mobility surveys.

¹⁰Bicycles and motos account for the remaining 1% of commutes. For comparison, the average commute in US cities has increased from 21 minutes in 1980 to 26 minutes in 2015.

city's bus fleet operated illegally (Cracknell 2003).¹¹ Disregard of bus stops promoted boarding and alighting along curbs, further reducing traffic flows.

The result was that while the crowding of Bogotá's streets slowed traffic overall, buses were much slower than cars. Table 2 compares speeds between buses and cars in 1995. Column (1) shows that commutes by car were around 35% faster than by bus. This is robust to controlling for differences in trip composition with trip origin-destination fixed effects in column (2). However, columns (3) and (4) show that low-skill Bogotanos were about 29% more likely to use buses than cars. The burden of slow public transit therefore fell disproportionately on the city's low-skill population.

3.2 TransMilenio: The World's Most Used BRT System

Background 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.¹² Phases 2 and 3 added an additional 70km in 2006 and 2011, 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 2013).¹³ Its average operational speed of 26.2kmh reported during phase one is on par with that of the New York subway (Cracknell 2003; Johnson 2010), 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.¹⁴ 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 while the government sets fares close to those offered by existing buses.¹⁵

There are two main reasons why BRT provides an attractive alternative to subways in rapidly growing

¹¹The Department of Mobility estimated the number to be more than double the amount actually required. A typical practice through which bus companies avoided government controls was duplication of license plates and vehicle documentation.

¹²While the anticipation of a system may predate its inauguration, TransMilenio went from a "general idea" to implementation in only 35 months (Hidalgo and Graftieux 2005). 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. The policy did not have the intended consequence of reduced car use: rising pollutants suggest increased purchases of old vehicles (Lawell et. al. 2016) and I report a mild increase in car ownership over the period in the appendix. This matches the experience of Mexico City (Davis 2008).

¹³A map of each system component and their opening date is provided in the appendix. 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.

¹⁴As shown in the appendix, concrete barriers separate TransMilenio from other lanes and have helped the city to achieve essentially complete compliance.

¹⁵For 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.

cities. First, it delivers similar reductions in commuting times at a fraction of the cost: the average per kilometer construction cost is one-tenth of rail (Menckhoff 2005). Second, BRT is much faster to construct. An illustrative comparison is that of TransMilenio and Metro de Medellín, the subway system in Colombia's second largest city. While both achieve similar system speeds, Medellín's metro cost eleven times as much as TransMilenio and took twelve years from announcement to opening.¹⁶

While BRT is not without drawbacks, these features have led to systems being built in more than 200 cities, the vast majority constructed over the past 15 years in Latin America and Asia (BRT Data 2017).

Route Selection and System Rollout The corridors built during the first phase of the system were consistently mentioned in 30 years of transportation studies as first-priority for mass transit (Cracknell 2003). The city conducted a planning study to reconfirm these suggested routes and identify new ones based on (i) current and future demand level and (ii) expected capital costs. The result was a plan that aimed to connect 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.¹⁷

Three features make TransMilenio an attractive context for empirical analysis. First, 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. Second, lines were placed along wide arterial roads that were cheaper to convert which were determined by the city's historical evolution. I leverage both in constructing instruments for the system's layout. Third, TransMilenio 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, combined with the staggered opening of lines across three phases, provide sources of time series variation I use in my analysis.

Finally, one central criticism of TransMilenio was its singular focus on improving urban mobility without coordinated changes in land use regulation (Bocajero et. al. 2013). As a result, I show in the appendix that housing supply did not respond to the system's construction. An integrated land use and transit policy tailored towards increasing housing densities near stations promotes a more efficient urban structure where many residents can 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.

¹⁶The difference in times from planning to opening between BRT and subways is not specific to Bogotá. While a comprehensive source on construction times is hard to find, stories of such instances are not. In India planning for subways in Delhi and Bangalore started eighteen and eight years before inauguration respectively, while the BRT in Ahmedabad took only 4 years. New York's second avenue subway line opened on January 1st 2017 having been originally proposed in 1919. In Bogotá, there were a total of ten attempts to introduce heavy rail between 1947 and 1997, thwarted by high capital costs and vested interests of the public transportation sector (Lleras 2003).

¹⁷See Cracknell (2003) for discussion. This was confirmed through inspection of satellite images (see appendix). Since this is not always possible in other cities (e.g. Jakarta, Delhi), an interesting extension would be to assume car and bus speeds fell along TransMilenio routes to assess the impacts in these contexts. That certain routes already contained median busways did not mean that there was efficient bus transit available along them (e.g. Avenida Caracas). Within a few years of their opening in 1990 the "the scheme became anarchic as, for example, (i) buses competed for passengers and this, together with little effective stop regulations, resulted in bus stop congestion and hazardous operating conditions, (ii) buses without a license to operate on Av. Caracas were attracted to the busway seeking passengers" (Cracknell 2003).

Trip Characteristics In the appendix, I provide additional details on the way in which TransMilenio is used and its effects on other modes which I summarize here. First, TransMilenio is a quantitatively important mode of transit that is more likely to be for longer trips compared to other modes. Second, TransMilenio provides an improvement in door-to-door speeds of around 17% over existing buses, but remains around 8.1% slower than cars. Third, the system is more likely to be used for commutes to work rather than leisure trips compared to other modes, motivating the focus on access to jobs in this paper. Fourth, TransMilenio use appears to have come primarily from substitution away from buses. Fifth, conditional on car ownership the rich and poor are equally likely to use TransMilenio, consistent with the similar fares as traditional buses.

Impact on Congestion BRT may affect equilibrium speeds through impacts on travel mode and route choices, and the number of lanes available for other traffic. In the case of Bogotá, the number of lanes available for other traffic was left unchanged: one might expect TransMilenio to have reduced congestion faced by cars and other buses. In the appendix, I show that there were in fact no significant changes in car and bus speeds along routes whose least cost paths lay along TransMilenio lines relative to other routes. This could be explained by substitution across modes and routes arbitraging any initial speed differences caused by the BRT, or a small elasticity of driving speeds to vehicles volumes at high levels of rush hour traffic.¹⁸ While incorporating congestion in the model would be an interesting extension, these moments suggest that my abstraction from the effects of TransMilenio on other mode speeds appears a reasonable approximation to reality.¹⁹

4 A Quantitative Model of a City with Heterogeneous Skills

This section presents a general equilibrium model of a city. There are a large number of discrete locations $i \in I$ that differ in their commute times to every other location, their housing floorspace as well as their amenities and productivities.²⁰ High- and low-skill workers decide where to live, whether to own a car, where to work, and which mode of transit to use to commute. Public transit is available to everyone to commute between home and work, but only those willing to pay to own a car have the option to drive. Firms from multiple industries are located across the city and produce using labor and commercial floorspace, and each industry differs in its demand for skills: for example, hotels and restaurants demand more low-skilled workers while

¹⁸The former is consistent with the “fundamental law of road congestion”: Duranton and Turner (2012) find that vehicle-kilometers travelled (VKT) increase one for one with roadway lane kilometers, and find no evidence that the provision of public transportation affects VKT. Akbar and Duranton (2017) estimate congestion in Bogotá and find the elasticity of speed with respect to the number of travelers is only 0.06 during peak hours, while Akbar et. al. (2017) find that only 15% of differences in driving speeds in Indian cities are due to congestion.

¹⁹Recent work has begun to incorporate congestion in trade models (Fajgelbaum and Schaal 2017; Allen and Arkolakis 2019). While allowing for congestion could have welfare implications even if speeds are unchanged (due to movements in traffic quantities), I am likely to be underestimating the welfare effect in its presence. When I remove the system, I keep the times on other buses and cars unchanged and measure the change in welfare from the observed equilibrium with TransMilenio. In a world with congestion, moving passengers from TransMilenio to other modes would slow speeds, in turn increasing the welfare change. Note my empirical results only speak to relative changes in speeds and are silent on the overall effect of TransMilenio on speeds. In the data, the aggregate speeds on cars and (non-TransMilenio) buses is uncorrelated with the system’s ridership: speeds fall significantly between 1995 and 2005 (a period of significant population growth of over 29%) while stabilizing between 2005 and 2015. This highlights the role of external aggregate shocks, such as urbanization lead by the country’s civil war, that motivates the more local analysis pursued in this paper.

²⁰The choice to keep the total supply of floorspace fixed is motivated by the result that this is mostly unaffected by TransMilenio as documented in the appendix. In Section 10, I explore the impact of allowing floorspace supply to respond to the system.

financial services require more high-skilled labor. Demand for skill therefore varies across the city based on the productivity of each industry in each location. Landowners choose how to allocate the fixed amount of floorspace across residential or commercial use. In equilibrium, the price of floorspace, the share allocated to each use and wages adjust to clear land and labor markets. This setup differs from recent quantitative urban models (e.g. Ahlfeldt et. al. 2015) by incorporating multiple skill groups of workers, commute modes and industries, where workers have non-homothetic demand for cars and residential amenities.

Despite the interactions between labor and land markets across thousands of locations that occur through the city’s commuting network, a single measure—CMA—summarizes the effect of the entire network on a location. This will be integral to my empirical and structural analysis in Sections 7 and 8 respectively.

4.1 Workers

The city is populated different worker skill groups indexed by $g \in G = \{L, H\}$ with a fixed population \bar{L}_g .²¹ A worker ω in group g chooses a location i in which to live, a location j in which to work, and whether or not to own a car denoted by $a \in \{0, 1\}$. Individuals derive utility from consumption of a freely traded numeraire good ($C_i(\omega)$); consumption of residential floorspace ($H_{Ri}(\omega)$); an amenity reflecting common components of how members of that group enjoy living in i under car ownership a (u_{iag}); and have a disutility from commuting that reduces their productivity at work ($d_{ija} \geq 1$). Workers are heterogeneous in their match-productivity with firms in each location ($\epsilon_j(\omega)$) and their preference for each residence-car ownership pair ($\nu_{ia}(\omega)$).

Commute costs differ by car ownership because car owners can choose between commuting by car or public transit (such as walking, bus or TransMilenio), whereas individuals without cars can only choose between public modes. Cars also provide an amenity benefit capturing the potential for improved leisure benefits, but come at a fixed cost of ownership $p_a > 0$.²²

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

$$\begin{aligned} \max_{C_i(\omega), H_{Ri}(\omega)} \quad & u_{iag} C_i(\omega)^\beta (H_{Ri}(\omega) - \bar{h})^{1-\beta} \nu_{ia}(\omega) \\ \text{subject to} \quad & C_i(\omega) + r_{Ri} H_{Ri}(\omega) + p_a a = \frac{w_{jg} \epsilon_j(\omega)}{d_{ija}} \end{aligned}$$

Solving for the optimal demand for housing and consumption good yields the following expression for indirect utility

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

where the iceberg commute cost $d_{ija} = \exp(\kappa t_{ija})$ increases with the time t_{ija} it takes to commute between i and j under car ownership a . The parameter $\kappa > 0$ controls the size of these commute costs.²³

²¹This is the “closed city” assumption: in Section 9 I allow for perfect mobility between the city and the rest of the country.

²²In the appendix, I outline a third stage mode choice problem in which individuals decide how to commute between home and work conditional on their car ownership decision. Car owners can choose between cars and public modes (walk, bus, TransMilenio) while non-car owners may only use public transportation. The result is that car owners face different average commute times for each trip.

²³In the appendix I show the semi-log gravity equation implied by this specification of commute costs fits the commuting data well,

In contrast to models with homothetic preferences, the fixed nature of expenditures on cars and housing allows me to match the Engel curves I document for car ownership and housing expenditure,²⁴ and drives sorting of workers over car ownership and residential neighborhoods by income. When cars are quicker than public modes of transit, the rich are more willing to pay the fixed cost since their value of time is higher. Similarly, the fixed expenditure on subsistence housing means that the poor spend a greater share of income on housing and are attracted to low amenity neighborhoods where it is cheaper.

Timing Workers first choose where to live and whether or not to own a car, and then choose where to work. I solve their problem by backward induction. This structure simplifies the model's algebra; in Section 9 I show the results are qualitatively similar if all choices are made simultaneously.

4.1.1 Employment Decisions

Having chosen where to live i and whether or not to own a car a , individuals draw a vector of match-productivities with firms in locations across the city from a multivariate Frechet distribution²⁵

$$F_g(\epsilon_1, \dots, \epsilon_J) = \exp \left(- \left[\sum_j \tilde{T}_g \epsilon_j^{-\frac{\tilde{\theta}_g}{1-\rho_g}} \right]^{1-\rho_g} \right).$$

The parameter $\tilde{\theta}_g$ measures the dispersion of productivities for type- g workers (comparative advantage), with a higher $\tilde{\theta}_g$ corresponding to a smaller dispersion, while the parameter ρ_g determines the correlation of an individual's talent across locations (absolute advantage). If $\rho_g = 1$ then draws are perfectly correlated within individuals while if $\rho_g = 0$ then they are perfectly uncorrelated. The scalar \tilde{T}_g controls the overall level of productivities for workers in a particular group.

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

$$\pi_{j|ia} = \frac{(w_{jg}/d_{ija})^{\theta_g}}{\sum_s (w_{sg}/d_{isa})^{\theta_g}} \equiv \frac{(w_{jg}/d_{ija})^{\theta_g}}{\Phi_{Riag}} \quad (2)$$

where $\theta_g \equiv \tilde{\theta}_g/(1 - \rho_g)$ reflects the relative strength of comparative advantage.

Individuals are more likely to commute to a location when it pays a high wage net of commute costs (the numerator) relative to those in all other locations (the denominator). The sensitivity of employment decisions to commute costs is governed by the dispersion of productivity. When workers have similar matches with firms in different locations (high θ_g), then commuting decisions are more sensitive to commute costs. Differences in

in line with the consensus in the literature (Fortheringham and O'Kelly 1989). I assume commute costs reduce labor supply rather than overall utility, since this simplifies the gravity equation for commute flows. In Section 9, I show the welfare effect is similar if commute costs reduce utility directly.

²⁴See the appendix for both figures and explanations of their construction.

²⁵This can be microfounded by a process of undirected job search where workers and firms meet according to a poisson process with match-productivity learned after each meeting.

productivity heterogeneity across skill groups will important in determining the incidence of commute costs, since it controls the extent to which individuals are willing to bear high commute costs to work in a location.

Resident Commuter Market Access Expected income prior to drawing the vector of match productivities is directly related to the denominator in (2) through

$$\bar{y}_{iag} = T_g \Phi_{Riag}^{1/\theta_g}, \quad (3)$$

where T_g is a transformation of the location parameter of the Frechet distribution.²⁶ I refer to Φ_{Riag} as *Resident Commuter Market Access* (RCMA). This summarizes the effect of the entire commuting network on the supply of residents to a location: it rises when a location is close (in terms of commute costs) to well-paid jobs.

4.1.2 Residential Location and Car Ownership Decisions

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

$$L_{Riag} = \lambda_U \left(u_{iag} \tilde{y}_{iag} r_{Ri}^{\beta-1} \right)^{\eta_g} \quad (4)$$

where $\tilde{y}_{iag} \equiv \bar{y}_{iag} - p_a a - r_{Ri} \bar{h}$ is income net of fixed expenditures and λ_U is an equilibrium constant. Workers are attracted to locations with high amenities, high net incomes and low house prices, with an elasticity determined by the dispersion of their idiosyncratic preferences η_g . The entire transit network only matters for individuals' residential choices in so far as it affects RCMA through expected income in (3).²⁷

4.1.3 Aggregation

Firm Commuter Market Access and Labor Supply Using the commuting probabilities (2), the supply of workers to any location is found by summing over the number of residents who commute there $L_{Fjg} = \sum_{i,a} \pi_{j|iag} L_{Riag}$. This implies

$$L_{Fjg} = w_{jg}^{\theta_g} \Phi_{Fjg} \quad (5)$$

²⁶The constants in this section are given by $T_g \equiv \gamma_{\theta,g} \tilde{T}_g^{1/\theta_g}$, $\gamma_{\theta,g} = \Gamma\left(1 - \frac{1}{\theta_g(1-\rho_g)}\right)$, $\lambda_U = \bar{L}_g (\gamma_{\eta,g} / \bar{U}_g)^{\eta_g}$ and $\gamma_{\eta,g} = \Gamma\left(1 - \frac{1}{\eta_g}\right)$ where $\Gamma(\cdot)$ is the gamma function and \bar{U}_g is average utility for group- g individuals. Expected utility prior to learning match productivities is $U_{iag\omega} = u_{iag} (\bar{y}_{iag} - p_a a - r_{Ri} \bar{h}) r_{Ri}^{\beta-1} \nu_{ia\omega}$.

²⁷Locations are populated by members of group g only if they are desirable ($\bar{u}_{iag} > 0$) and affordable ($\bar{y}_{iag} - p_a a - r_{Ri} \bar{h} > 0$). The expression for residential populations (4) therefore applies only for "active" locations $\mathcal{A}_{Rg} = \left\{ (i, a) : \bar{u}_{iag} > 0, r_{Ri} < (\Phi_{Riag}^{1/\theta_g} - p_a a) / \bar{h} \right\}$ that are both desirable and affordable and is zero otherwise. For clarity, I omit this additional notation in the text. Note that while I ensure workers have positive income on average in each location, inspection of (1) shows that since $\epsilon_j(\omega)$ affects income there will be a small measure of workers who ex-post cannot afford the subsistence requirement. This would be resolved by having a preference rather than productivity shock for each workplace location, or by having a government who provides a lump sum transfer to ensure individuals can afford to live in each location. In Section 9, I show the welfare effect is similar under preference shocks so this is unlikely to affect the results.

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

Labor supply in the model takes a log-linear form that depends on two forces. First, more workers commute to destinations paying higher wages. Second, firms attract workers when they have better access to them through the commuting network, captured through the term Φ_{Fjg} . This is because individuals care about wages net of commute costs. I define the term Φ_{Fjg} as a location's *Firm Commuter Market Access* (FCMA). It summarizes the effect of the entire commuting network for firms in a location through its effect on labor supply. Total effective labor supply to location is given by $\tilde{L}_{Fjg} = \bar{\epsilon}_{jg} L_{Fjg}$, where $\bar{\epsilon}_{jg}$ is the average productivity of type- g workers who decide to work in j .²⁸

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

$$\bar{U}_g = \gamma_{\eta,g} \left[\sum_{i,a} \left(u_{iag} \tilde{y}_{iag} r_{Ri}^{\beta-1} \right)^{\eta_g} \right]^{1/\eta_g} \quad (6)$$

4.2 Firms

Technology There are $s \in \{1, \dots, S\}$ industries which produce varieties differentiated by location in the city under perfect competition. Output is freely traded, and consumers have CES preferences over each variety with elasticity of substitution $\sigma_D > 1$.²⁹ Firms produce using a Cobb-Douglas technology over labor and commercial floorspace

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

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

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

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

Factor Demand Perfect competition implies that the price of each variety is equal to its marginal cost

²⁸In particular, $\bar{\epsilon}_{jg} = T_g \sum_{i,a} \frac{\pi_{j|ia}^{-1/\theta_g}}{d_{ija}} \frac{\pi_{j|ia} L_{Riag}}{\sum_{r,o} \pi_{j|rog} L_{Rrog}}$.

²⁹This is the numeraire good introduced in the consumer's problem. While the assumption of representative firms with a fixed location seems restrictive, in the next section I show this production technology is isomorphic to more realistic setups.

$p_{js} = W_{js}^{\alpha_s} r_{Fj}^{1-\alpha_s} / A_{js}$, where r_{Fj} is the price of commercial floorspace in j and

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

is the cost of labor for firms of industry s in location j . Intuitively, labor costs differ by industries due to their differential skill requirements. Solving the firm's cost minimization problem and letting X_{js} denote firm sales, the demand for labor and commercial floorspace is³⁰

$$\tilde{L}_{Fjgs} = \left(\frac{w_{jg}}{\alpha_{sg} W_{js}} \right)^{-\sigma} N_{js} \quad (7)$$

$$H_{Fjs} = (1 - \alpha_s) \frac{X_{js}}{r_{Fj}}. \quad (8)$$

4.3 Floorspace

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

$$r_{Ri} = (1 - \beta) \frac{E_i}{H_{Ri} - \beta \bar{h} L_{Ri}} \quad (9)$$

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

$$r_{Fj} = \frac{\sum_s (1 - \alpha_s) \left(W_{js}^{\alpha_s} r_{Fj}^{1-\alpha_s} / A_{js} \right)^{1-\varsigma} X}{H_{Fj}}. \quad (10)$$

Floorspace Use Allocation Landowners choose the fraction ϑ_i of floorspace allocated to residential use to maximize profits. They receive r_{Ri} per unit of floorspace allocated to residential use, but land use regulations limit the return to each unit allocated to commercial use to $(1 - \tau_i) r_{Fi}$. Landowners allocate floorspace to its most profitable use so that

$$\begin{aligned} \vartheta_i &= 1 \text{ if } r_{Ri} > (1 - \tau_i) r_{Fi} \\ (1 - \tau_i) r_{Fi} &= r_{Ri} \quad \forall \{i : \vartheta_i \in (0, 1)\} \\ \vartheta_i &= 0 \text{ if } (1 - \tau_i) r_{Fi} > r_{Ri} \end{aligned} \quad (11)$$

³⁰From CES demand $X_{js} = p_{js}^{1-\sigma} X$ where $X = \sum_i \beta (E_i - \bar{h} r_{Ri} L_{Ri})$ is total spending on goods in the city and $E_i = \sum_{g,a} (\bar{y}_{iag} - p_a \alpha) L_{Riag}$ is total spending on goods and housing from residents in i .

4.4 Externalities

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

$$A_{js} = \bar{A}_{js} \left(\tilde{L}_{Fj} / T_j \right)^{\mu_A}, \quad (12)$$

where $\tilde{L}_{Fj} = \sum_s \tilde{L}_{Fjs}$ is the total effective labor supplied to that location and T_j is the total units of land. The strength of agglomeration externalities is governed by the parameter μ_A .³¹

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

$$u_{iag} = \bar{u}_{iag} \left(L_{RiH} / L_{Ri} \right)^{\mu_{U,g}}. \quad (13)$$

In contrast to existing urban models (e.g. Ahlfeldt et. al. 2015), endogenous amenities depend on the composition of residents across skill groups rather than the total density of residents. This seems especially applicable in developing country cities that lack strong public goods provision. In Bogotá, where crime is a significant problem, the rich often pay for private security around their buildings which increases the sense of safety in those areas. This externality provides an additional force towards residential segregation, since the high-skilled are more willing to pay to live in high-amenity neighborhoods and by doing so increase the amenities even more. While sorting could be driven by the subsistence housing requirement alone, I allow the strength of residential externalities $\mu_{U,g}$ to differ across groups and let the data speak to the relative strength of these forces towards residential segregation in estimation.³²

4.5 Equilibrium

I now define general equilibrium in the city.

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

1. **Labor Market Clearing** The supply of labor by individuals (5) is consistent with demand for labor by firms (7),

³¹I do not allow for spillovers across locations since spatial units in my analysis are census tracts given the evidence on highly localized spatial spillovers (Rossi-Hansberg et. al. 2010; Ahlfeldt et. al. 2015). However, earlier versions of the paper show how the regression approach can incorporate such spillovers.

³²Apart from improving accessibility, TransMilenio may have direct impacts on locations productivities \bar{A}_{js} and amenities \bar{u}_{iag} (e.g. through street improvements, effects on crime or pollution). In Section 8 I allow for this possibility but find no economically significant effects, motivating their exclusion from the model.

2. **Floorspace Market Clearing** The market for residential floorspace clears (9) and its price is consistent with residential populations (4), the market for commercial floorspace clears (10) and floorspace shares are consistent with land owner optimality (11),
3. **Closed City** Populations add up to the city total, i.e. $\bar{L}_g = \sum_{i,a} L_{Riag} \forall g$.

With this definition in hand, I now characterize the equilibrium.

Proposition 1. *An equilibrium exists in this city. Moreover, in a special case of the model with one group of workers, firms and commute modes and no non-homotheticities ($\bar{h} = p_a = 0$) and a fixed allocation of floorspace, a sufficient condition for the equilibrium to be unique is that*

$$\begin{aligned}\mu_A &\leq 1 - \alpha + \frac{\sigma + \theta - 1}{(\sigma - 1)(\theta - 1)} \\ \mu_U &\leq \frac{1 + \eta(1 - \beta)}{\eta} - \frac{\beta}{\theta - 1} \\ \beta(\sigma - 1)\mu_A + \sigma\mu_U &\leq \frac{\sigma}{\eta} + \sigma(1 - \beta) + \beta(1 + (\sigma - 1)(1 - \alpha))\end{aligned}$$

Despite adding multiple groups of workers, industries and transit modes, and two fixed expenditures that potentially influence the set of locations inhabited by each group, the first part of the proposition ensures an equilibrium still exists. The second part provides intuition for equilibrium multiplicity by showing that in the simplified model, the equilibrium will be unique only when spillovers are sufficiently weak. In the presence of strong spillovers, fundamental productivities and amenities become less important and different urban configurations can be supported as equilibria. While multiplicity does not pose a problem for my estimation strategy which only requires that two equilibria be observed, I address it through my equilibrium selection rule when solving for counterfactual equilibria.

5 Using The Model To Guide Empirical Work

In this section, I show that in a special case the model's equilibrium has a reduced form representation in which outcomes such as population, employment and floorspace prices can be written as log-linear functions of CMA. This framework applies to a wide class of models featuring a gravity equation for commute flows and thus is robust to alternative modeling assumptions.

I use this regression framework primarily as a model validation exercise to examine whether the city responds to the change in infrastructure as predicted by a class of gravity of models. The full model with multiple layers of heterogeneity does not have an exact log-linear representation; I structurally estimate it using similar moment conditions in Section 8.³³ I also use the reduced form elasticities directly in the first order welfare

³³One concern is how valid this is a model validation exercise, since it abstracts from the heterogeneity I claim to be in the data generating process. Log-linearizing the full equilibrium equations would deliver similar specifications. For example, the change in residential house prices would depend on a weighted average of the RCMA shocks for each skill group and car ownership cell, with heterogenous coefficients reflecting a location's composition. One interpretation is then that the simplified regression is approximating the average effect from the full model. In the appendix, I conduct a Monte Carlo exercise in which I run the same regression specifications on data simulated from my full model and show that the non-parametric relationships between outcomes and CMA in the simplified regression model are indeed log-linear.

approximations in Section 9.

Commuter Market Access: Measurement and Intuition Consider a simplification of the model with one group of workers, firms and transit modes and a fixed allocation of residential and commercial floorspace. Using the labor supply curve (5) to substitute for wages in the expression for RCMA in (2), the CMA terms can be expressed (to scale) as the solution to the following system of equations

$$\Phi_{Ri} = \sum_j d_{ij}^{-\theta} \frac{L_{Fj}}{\Phi_{Fj}} \quad (14)$$

$$\Phi_{Fj} = \sum_i d_{ij}^{-\theta} \frac{L_{Ri}}{\Phi_{Ri}} \quad (15)$$

RCMA reflects access to well-paid jobs. It is greater when a location is close (in terms of commute costs) to other locations with high employment, particularly so when these other locations have low 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 have low access to jobs (lowering the wage individuals are willing to work there for). I show below that the solution to this system of equations exists and is unique, so that the market access measures are easily computed using data on population, employment and commute costs as well as a value for the commuting elasticity.

Regression Framework In the appendix, I show that in this simplified model the equilibrium can be written as the following system

$$A \Delta \ln Y_i = B_R \Delta \ln \Phi_{Ri} + B_F \Delta \ln \tilde{\Phi}_{Fi} + e_i$$

where $\Delta \ln Y = \left[\Delta \ln L_{Ri} \quad \Delta \ln r_{Ri} \quad \Delta \ln r_{Fi} \quad \Delta \ln \tilde{L}_{Fi} \right]'$ is a vector of log changes in endogenous variables, A is a matrix of coefficients reflecting the interdependence between endogenous variables, B_R and B_F are vectors of coefficients controlling the direct effects of changes in market access on outcomes, e is a vector of structural residuals containing changes in fundamentals and $\tilde{\Phi}_{Fi} = \sum_i d_{ij}^{-(\theta-1)} \frac{L_{Ri}}{\Phi_{Ri}^{\frac{\theta}{\theta-1}}}$ is adjusted firm commuter access capturing access to effective units of labor.

Premultiplying by the coefficient matrix A yields the reduced form

$$\Delta \ln Y_i = A^{-1} B_R \Delta \ln \Phi_{Ri} + A^{-1} B_F \Delta \ln \tilde{\Phi}_{Fi} + A^{-1} e_i \quad (16)$$

The coefficient vectors $A^{-1} B_F$ and $A^{-1} B_R$ have a block structure with zeros in the first and last two entries respectively, so that each row collapses to a univariate specification in which residential (firm) outcomes depend only on residential (firm) commuter market access. This reduced form reflects the total change of outcomes in response to change in market access, comprised of both the direct effect (in the B_R and B_F coefficient vectors) and the indirect effect (in A^{-1}) as the response filters through labor and land markets.

The following proposition shows that this reduced form representation and the ability to retrieve measures of market access using only the gravity equation for commuting of commuters across the city is shared by a wide class of gravity models. For brevity, a complete formal statement of the proposition and its proof are provided in the appendix.

Proposition 2. (i) Measuring CMA *In a gravity commuting model with commute flows $L_{ij} = \gamma_i \delta_j \kappa_{ij}$ where $\gamma_i, \delta_j > 0$ are endogenous, the supply of residents and workers is log-linear and given by $L_{Ri} = \gamma_i \Phi_{Ri}$ and $L_{Fi} = \delta_i \Phi_{Fi}$, where Φ_{Ri}, Φ_{Fi} are uniquely determined (to scale) by data $\{L_{Ri}, L_{Fi}\}$ and parameters $\{\kappa_{ij}\}$.*

(ii) Isomorphisms *In a gravity commuting model with log-linear demand for residents and labor $\tilde{L}_{Fj} = A_j \delta_j^\alpha$ and $L_{Ri} = B_i \gamma_i^\beta \Phi_{Ri}^\gamma$ where $A_i, B_i > 0$ are exogenous constants and the supply of labor (potentially different from the number of workers) is given by $\tilde{L}_{Fj} = \delta_j^\delta \tilde{\Phi}_{Fj}$, where $\tilde{\Phi}_{Fj} = \sum_i \frac{L_{Ri}}{\Phi_{Ri}} \kappa_{ij} \Phi_{Ri}^\epsilon$, in equilibrium residence and effective employment can be expressed as log-linear functions of $\Phi_{Ri}, \tilde{\Phi}_{Fi}$ and constants A_i, B_i . The equilibrium always exists and is unique when $|\epsilon(\beta - 1) - \gamma| \leq |\beta - 1| |\alpha - 1|$.*

The gravity equation for commuting that determines the supply side of the model enjoys wide empirical support and is used in the vast majority of recent quantitative urban models.³⁴ The first part of the proposition shows that unique values of market access can be computed using data on the location of residence and employment, as well as a parameterization of commute costs, using only the supply side of the model through the gravity equation for commute flows. The second part shows that for a class of models with log-linear demand for residents and labor, equilibrium population and employment can be written as log-linear functions of CMA. In the appendix, I show that beyond my own model this framework accommodates iso-elastic housing supply, alternative production technologies (e.g. Eaton and Kortum 2002, and individual entrepreneurs who choose where to produce across the city) and worker preferences (such as utility over leisure). The regression framework I take to the data is therefore robust to a host of alternative modeling assumptions.³⁵

Relation to Market Access Literature In the economic geography literature, individuals live and work in the same location and goods trade is subject to trade costs. A class of models contain measures of market access reflecting the ease for consumers (firms) to buy (sell) goods from (to) other locations. Since this is silent on the impact of commuting infrastructure, I consider a different class of urban models where goods trade is costless but individuals can live and work in separate locations. These models contains measures of accessibility of residents (firms) to jobs (workers). The similarity is natural given that both rely on a gravity equation to model the flows of goods or factors across space.

CMA can also be recovered from observable data using less model structure than is typically used in economic geography settings. The only structure used to derive the system (14) and (15) is the gravity equation for commuting. In economic geography settings, additional assumptions such as symmetric trade costs, balanced trade and goods market clearing are often made to recover market access measures from the data. In fact, one

³⁴See McDonald and McMillen (2010) for a review of the evidence in support of gravity in commute flows; these include all the quantitative models in Section 2.

³⁵In the appendix, I also show the framework can be used when workers have preference rather than productivity shocks over employment locations, so that there is no difference between effective labor and the number of workers. Most models impose additional restrictions between $\alpha, \beta, \gamma, \delta$ and ϵ which reduces the number of parameters one needs to know (see the appendix for examples).

can show that it is precisely the absence of balanced trade in commuters that deliver separate notions of resident and firm CMA in this paper. This distinction is important given that changes in firm and resident CMA capture very different sources of variation.³⁶

6 Data

In this section I provide an overview of the primary datasets used in the analysis. Additional details are provided in the data appendix.

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 2005.³⁷ These are contained within larger spatial units including 19 localities and 113 planning zones (UPZs).

My primary source of population data is the Department of Statistics’ (DANE) General Census of 1993 and 2005. This provides the residential population of each block by education level. I define college-educated individuals as those with some post-secondary education defined by their last complete year of study. In 2015, DANE provides population totals at the UPZ. I combine this with the share of college-educated workers in each UPZ in the GEIH survey in that year (described below) to construct population by skill group. This allows me to compute separate growth rates of college and non-college residents between 2005 and 2015 within each UPZ. I then calculate 2015 census tract population by skill group by inflating the 2005 totals by these growth rates. Details are provided in the appendix.

I use two sources of data on employment. 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 this contains the location, industry and employment of each establishment, but in 2000 establishment size is not provided. While I tend to use the census and CCB datasets separately, a concern is that the spatial distribution of registered employment may be different from that of total employment. In the appendix, I show that the employment and establishment densities in both years of the CCB data are highly correlated with that from the 2005 census. Importantly, coverage is even across rich and poor neighborhoods, suggesting both that the CCB data is representative of overall employment. Since I rely on establishment counts to proxy for employment in the CCB data, I also show that establishment count and employment densities are highly correlated in years where both are available.

³⁶Balanced trade and symmetric trade costs imply that firm and consumer market access collapse to a single measure. While these measures are highly correlated in the economic geography setting, it makes less sense in an urban setting where balanced trade in commuters in my setting would require the number of workers in a location to equal the number of residents, which is counterfactual. The amount of additional model structure these papers impose is inversely related to the granularity of the data. Redding and Venables (2004) impose none of these additional restrictions but have much stronger data requirements: they need to observe trade flows between each geographic unit in their data. Bartelme (2015) only requires symmetric trade costs and balanced trade, while Donaldson and Hornbeck (2015) also impose goods market clearing since they only observe population rather than expenditure in each location. In contrast, I use only the gravity equation determining the supply of commuters to recover the CMA measures directly from data residential population and employment, which are typically available in cities from censuses and new sources such as cellphone metadata records.

³⁷Number reported is for tracts with positive population. Almost all tracts (2,768) have positive population in 2005. For comparison, tracts in Bogotá are about 60% smaller than those in New York City which had an average of 4,067 residents in the 2010 census.

Housing market data between 2000 and 2012 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 with the result that the dataset covers 98.6% of the city's more than 2 million properties (Ruiz and Vallejo 2015).³⁸ 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. As described in the appendix, 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 2015).³⁹ As a result, I show 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 demographic 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 details on employment such as income, hours worked and industry of occupation across primary and secondary jobs. I was able to access 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.⁴⁰ 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.

Finally, I measure the distance of tracts to various spatial features provided by the city. I also use a land

³⁸I confirmed this high coverage 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 2015).

³⁹Surveyors are sent out to update the characteristics of each property every couple of years. Since the primary data informative about prices is not necessarily updated each year, I focus on long-differences in my analysis.

⁴⁰For example, the TransMilenio network is the union of pedestrian paths, trunk lines and feeder routes; the latter two can only be entered at stations. As described in detail in the appendix, one mode may have different speeds depending on the part of the network. For example, cars have different speeds on primary, secondary and tertiary roads.

use map of the city in 1980 provided by the US Defense Mapping Agency and a Tramway map from Morrison (2007).

7 Empirical Analysis

In this section, I use the log-linear relationships between endogenous outcomes and CMA derived in Section 5 to empirically assess the effect of TransMilenio on land and labor markets.

7.1 Approach and Identification

Consider the residential outcomes in the first two entries of the reduced form system (16)

$$\Delta \ln Y_{Rit} = \beta \Delta \ln \Phi_{Rit} + \alpha_\ell + \gamma' X_{it} + \epsilon_{Rit} \quad (17)$$

That is, I regress changes in (log) residential outcome Y_{Rit} in census tract i in year t on changes in (log) RCMA Φ_{Rit} , as well as a set of controls that contain census tract characteristics X_{it} and locality fixed effects α_ℓ to partially capture the unobservables in the structural residential. An equivalent specification holds for firm outcomes, which instead depend on FCMA. These CMA terms are defined by the system of equations in population, employment and commute costs $d_{ijt}^{-\theta}$ in (14) and (15). The elasticity β is identified from variation in CMA within census tracts over time, comparing tracts within a locality with similar observable characteristics which experienced different changes in market access.⁴¹

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.⁴² 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.⁴³

Challenges to Identification There are two key challenges to estimating the specification (17). First, changes in CMA contain population and employment in both periods. Productivity and amenity shocks (in the error term) that drive movements in residence and employment will therefore be mechanically correlated with changes in CMA. To isolate the variation in CMA due only to changing commute costs, I instrument for the change in CMA when population and employment are fixed at their initial values in the system (14) and (15).⁴⁴

⁴¹The reduced form elasticities are outcome-specific but I omit this additional notation here. For firm outcomes, I use the unadjusted firm commuter access instead of the adjusted term that reflects units of effective labor supplied. The results are qualitatively unchanged if I use the adjusted term (the measures have a 0.99 correlation).

⁴²In order to compute the market access terms, I require values for $d_{ij}^{-\theta} = \exp(-\theta\kappa t_{ij})$. The estimation of θ, κ is outlined in the next section; I measure θ and t_{ij} by averaging over skill group and car ownership values respectively weighting by population shares of each category. 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.

⁴³Firm CMA also increases toward the center-North of the city due to the high density of (low-skill) workers in the South.

⁴⁴I exclude the location itself when calculating its predicted change in CMA to address the possible correlation between initial residence and employment and unobserved shocks. In robustness checks, I extend this to exclude all tracts within 1.5km of a location.

Second, CMA growth may be correlated with the error if TransMilenio routes targeted neighborhoods with differential trends in productivities or amenities. For example, the government may have wanted to support growing neighborhoods or to stimulate lagging ones. I therefore construct two instruments for TransMilenio routes, which in turn imply two instruments for the change in CMA.⁴⁵ The first instrument takes as given the government's overall strategy of connecting portals at the edge of the city with the CBD as given, excludes those areas from the analysis, and constructs the routes that would have been built if the sole aim had been to minimize costs. I construct these routes by first digitizing a land use map of Bogotá in 1980 to measure the different types of land use on small pixels across the city (e.g. arterial roads, vacant, developed etc). Using engineering estimates for the cost to build BRT on different types of land use, this provides a construction cost raster for the city based on the share of land use in each pixel. This allows me to solve for the least-cost paths connecting portals with the CBD. 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 should predict the location of TransMilenio since these 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.⁴⁶ 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.

My 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. This is much harder to implement in distance-based specifications where there is likely to be little residual variation in distance to one particular least cost road conditional on distance to main roads. To provide further evidence in support of my identification assumption, I check the stability of IV point estimates as controls are added and test that both instruments yield similar coefficients. I also run a host of robustness checks described below.

⁴⁵The CMA instruments are constructed as follows. In ArcMap, I compute the commute times had the system been built along each instrument. Plugging these into (14) and (15) 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 defined as the difference between this predicted CMA under TransMilenio and its value in the initial period without the system. Similar approaches have been used in the trade and economic geography literature to predict the location of present day infrastructure based on historical instruments (Baum Snow 2007; Duranton and Turner 2012) or predicted construction costs (Faber 2015, Alder 2017).

⁴⁶These include 1933 population and distance to main roads in 1933. That I extend the tram lines to the edge of the city should also reduce concerns over direct effects of the tram on outcome growth, since much of the instrument was never built.

7.2 Results: Main Outcomes

Main Outcomes Table 3 presents the main results. In all specifications, only tracts further than 500m from a portal and the CBD are included in order to keep a constant sample across specifications.⁴⁷ Columns (1) and (2) report the results from the OLS regressions where the change in CMA is measured using both the change in commute costs as well as the change in population and employment. In most cases, the point estimates are slightly lower in column (2) (my preferred specification with full controls) due to the positive correlation of changes in market access with initial land market and demographic characteristics that caused treated locations to grow faster over the period.

Columns (3) and (4) run the baseline IV specification, which instrument for the total change in market access holding employment and population fixed at their initial levels. The point estimates tend to fall slightly, reflecting the positive mechanical correlation previously discussed.

Columns (5) and (6) instrument for the change in market access both by holding initial employment and population constant and computing the change in commute times had TransMilenio been built along the least-cost path instrument. For residential outcomes, the point estimates are larger than columns (3) and (4). While this could be (partially) due to measurement error, the difference suggests a negative correlation between TransMilenio placement and growth in unobserved amenities and productivities. This seems plausible, given that the system was built to serve areas of the city that had been growing during the 1990s and may therefore have slowed down during the 2000s as they became congested. Commercial outcomes are more noisy, but the overall pattern is that the IV estimates are slightly higher than the previous estimates. That the estimates are stable as additional controls are added provides additional evidence in support of the exclusion restriction.

Finally, columns (7) and (8) use both the tram and LCP instruments. The coefficients remain stable compared to using the LCP instrument alone, and in all but one case I fail to reject validity of the overidentification restrictions.⁴⁸

Heterogeneous Effects of Transit Figure 4 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.

Robustness In the appendix, I report a number of additional results which I summarize here.

First, I use less model-dependent measures of resident and firm CMA. These are commute-time weighted

⁴⁷One limitation of my data is that variables do not line up over time periods and each specification may therefore rely on changes over different periods. However, I always use changes in market access constructed between the two waves in question and measure CMA using the values for population and employment in each period. Population regressions using differences from 1993 to 2005 measure changes in market access induced by phase one (opened between 2000 and 2003, with 47% and 73% of the phase opened by 2000 and 2001 respectively). Land market and employment regressions using differences between 2000-2012 and 2000-2015 respectively measure changes in market access induced by phase one and two (opened between 2005-2006). Employment and population regressions are weighted by initial establishment counts and population respectively to increase precision, but in robustness checks I show the results also hold in unweighted regressions. I also restrict the sample to tracts within 3km of stations for main specifications to ensure the results are not being driven by changes in CMA in very distant tracts, but include all tracts in robustness checks.

⁴⁸Comparison of these reduced form estimates with those implied by the model's structural parameters is provided in Section 9.1.

sums of employment and residence respectively, and recall the “market potential” discussed by Harris (1954) and alluded to in the discussion of accessibility in Hansen (1959).⁴⁹ Second, I run falsification tests to check that changes in CMA induced by particular lines are not associated with growth in outcomes before they open. Third, I condition on distance to stations to show that the effects are driven by changes in market access rather than characteristics of stations (e.g. changes in foot traffic, pollution or complementary infrastructure). Fourth, I assess the response of variables to changes in market access to distant locations more than 1.5km away, reducing the potential for bias resulting from local unobservables. Both of these empirical approaches are not possible with a distance-based empirical approach. Fifth, I use alternative speeds to compute the commute times for each mode. Sixth, I vary the commute elasticity θ to 1.5 and 0.5 times its estimated value. Seventh, I include all census tracts in the analysis, rather than those within 3km of a station. Eighth, I run unweighted regressions for employment and population regressions which are weighted in the main specifications. Finally, 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. That my results are robust to these alternative specifications provides additional evidence in support of the causal effect of TransMilenio on urban outcomes through improvements in CMA.

Comparison with Distance Band-Based Predictions One key benefit of the CMA approach is that by capturing the specific geography of a city and the change in its transit network, estimates from one context are more likely to port to others than those based on distance to stations. Another advantage is the rich heterogeneity in treatment effects shown in Figure 1. In the appendix, I compare the predictions for residential house price growth in the CMA model with those from a distance-based regression of the change in floorspace prices on two dummies for being closer than 750m from a station and between 750-1500m from a station (relative to the omitted tracts between 1.5-3km away). The dissimilarity index for the predicted changes is 0.631, with appreciation over- (under-)predicted in the center (outskirts).⁵⁰

7.3 Results: Additional Outcomes

Commute Distance Table 4 examines whether TransMilenio led to changes in commuting distances. Column (1) shows that changes in market access caused by TransMilenio were indeed associated with greater probability of using the system in 2015, providing reassurance that the measure captures changes in commuting opportunities. Columns (2)-(4) run difference-in-difference specifications similar to (17) exploring how changes in market access affected commute distances within residential locations (UPZs) between 1995 and 2015. Throughout the OLS and IV specifications, improvements in CMA led to increases in commute distances, suggesting the system made employment in more distant locations more attractive. Finally, column (5) tests for heterogeneous effects across workers and finds the effect on commute distances is mildly greater for low-skill workers. This likely reflects both their greater reliance on public transit as well as a greater sensitivity

⁴⁹In particular, I define $RMP_i = \sum_j t_{ij}^{-1} L_{Fj}$ and $FMP_i = \sum_j t_{ij}^{-1} L_{Rj}$ as resident and firm market potential respectively. I also add each control variable incrementally to the OLS specification to show its impact on the CMA coefficient.

⁵⁰For two variables X_i, Y_i this is defined as $\frac{1}{2} \sum_i \left| \frac{X_i}{\sum_k X_k} - \frac{Y_i}{\sum_k Y_k} \right|$. It varies between zero and one, with zero indicating identical distributions across locations.

of commute flows to commute costs as shown in the next section.

College Share A key question surrounding the debate on the effects of public transit is whether it leads to a re-sorting of worker (skill) groups. In the US, investments in transit have typically been followed by reductions in the share of rich residents (e.g. Glaeser et. al. 2008) although there is evidence this effect varies across different types of neighborhoods (Heilman 2017). However, the evidence in developing countries is far sparser.

Table 5 explores how the share of college residents in a census tract responds to changes in RCMA. Column (1) shows that, on average, there was no significant effect on demographic composition. However, this may mask underlying heterogeneity: the model predicts that the high-skilled are more willing to pay for improved access to jobs in neighborhoods with high amenities.⁵¹ In columns (2) to (4), I test whether the response differed by tracts according to the college share of the surrounding neighborhood.⁵² The college share did increase in response to an increase in market access, but only in neighborhoods with an initially high college share. In other words, the high-skilled were only willing to pay for improved transit access in “nicer” neighborhoods, and would not trade off these benefits for the lower amenities in poorer locations in the South. In contrast, the low-skilled were more likely to move into poorer neighborhoods with a lower initial college share. Overall, this shows that TransMilenio increased residential segregation between the low- and high-skilled.

Wages Table 6 examines the impact of market access on wages reported by individuals across UPZs. I run a difference-in-difference specification similar to (17) to examine the effect of improved RCMA on log average hourly wages reported by full-time workers between 18 and 55. Column (1) shows a strong association between improved access to jobs and wages over the period. However, column (2) controls for the changing educational composition of workers and shows that about half of the relationship is explained by re-sorting of workers by skill. The result is qualitatively unchanged when using the IVs in columns (3) and (4). Finally, column (5) shows that the effect of RCMA on wages is greater for high-skilled individuals than for the low-skilled. While my cross-sectional data do not allow me to control for individual fixed effects, that wages rise even when controlling for changing worker characteristics supports the idea that CMA reflects accessibility to high-paid jobs. That the effect is greater for high-skill workers suggests they may benefit more from improved transit, a topic I return to in the quantitative section.⁵³

⁵¹ Assuming one mode of transit for the moment and the same wage for skill groups, log-linearizing the expression for residential populations (4) yields

$$\Delta \ln L_{Rig} \approx \eta_g \frac{\mu_{ig}^Y}{\theta_g} \Delta \ln \Phi_{Ri} - \eta_g (1 - \beta + \mu_{ig}^R) \Delta \ln r_{Ri} + \eta_g \Delta \ln u_{ig}$$

where $\mu_{ig}^Y = \frac{T_g \Phi_{Ri}^{1/\theta_g}}{T_g \Phi_{Ri}^{1/\theta_g} - r_{Ri} \bar{h}}$ and $\mu_{ig}^R = \frac{r_{Ri} \bar{h}}{T_g \Phi_{Ri}^{1/\theta_g} - r_{Ri} \bar{h}}$. Note μ_{ig}^R is greater for poor individuals in expensive neighborhoods. Thus poor, low-skilled workers are more sensitive to house price appreciation in expensive neighborhoods and are less willing-to-pay for improved CMA than the high-skilled.

⁵² I measure a tract’s surrounding college share using the share of college residents within a 1km disk around each tract centroid in 1993 (excluding the tract itself). I then define a high college dummy equal to one for tracts in the top two terciles of its distribution. The results are robust to using own-tract college share (subject to mechanical bias as a lagged dependent variable).

⁵³ Effects on other labor market outcomes are available upon request. I find a mild fall in hours worked and the probability an individual is employed at a small establishment with less than 5 workers, but the vast majority of these effects are driven by changing skill-composition of residents.

8 Structural Estimation

In this section I structurally estimate the full model from Section 4 which allows a quantitative assessment of the distributional effects of TransMilenio. I first describe how the model can be inverted to obtain the unobservable wages, amenities and productivities that rationalize the observed data as an equilibrium of the model. I then outline the procedure to estimate the model's parameters. Finally, I present the estimation results and model diagnostics.

8.1 Model Inversion

The model contains unobserved location characteristics, such as wages, productivities, amenities and land use wedges. While the presence of agglomeration forces allows for the possibility of multiple equilibria, I am able to recover unique values of composite productivities and amenities that rationalize the observed data as a model equilibrium.

There is a key difference in this process compared to recent quantitative urban models (e.g. Ahlfeldt et al. 2015). In those models, there is one group of workers. Combining data on residence and employment and leveraging the model structure provided by the gravity equation in commuting, it is straightforward to solve for the unique vector of wages that rationalizes the observed data. To replicate this in a model with multiple skill groups would require data on residence and employment by skill group. While data on residence by skill group are typically available from population censuses, I am unaware of similar datasets that provide employment by skill group within small spatial units within cities. This is where the model's multiple industries becomes useful. In the data I observe employment by industry. Intuitively, given the different skill-use intensities across industries, the relative employment by industries in a location should be informative about the relative employment across skill groups. The following proposition formalizes this intuition, and shows that a unique vector of group-specific wages can be recovered using data on residence by skill and employment by industry. Obtaining the remaining unobservables is straightforward.

Proposition 3. (i) *Wages* Given data on residence by skill group L_{Rig} , employment by industries L_{Fjs} , commute costs d_{ija} and car ownership shares $\pi_{a|ig}$ in addition to model parameters, there exists a unique vector of wages (to scale) that rationalizes the observed data as an equilibrium of the model.

(ii) *Remaining Unobservables* Given model parameters, wages and data $\{L_{Rig}, \pi_{a|ig}, L_{Fjs}, H_i, \vartheta_i, r_{Ri}, r_{Fi}\}$ there exists a unique vector of unobservables $\{u_{iag}, A_{js}, X_{js}, \tau_i\}$ (to scale) that rationalizes the observed data as an equilibrium of the model.

8.2 Parameter Estimation

My procedure to estimate the parameters of the model can be summarized as follows. First, I calibrate and estimate a subset of parameters without solving full model. Second, I solve for wages using parameters from the first step. Third, I estimate the remaining elasticities via GMM using moments similar to reduced form analysis. Fourth, with all parameters in hand, I invert the model to recover unobservables.

8.2.1 Parameters Calibrated to Exogenous Values

I calibrate $\{\sigma, \sigma_D, \alpha_s\}$ to existing values from the literature. I set the elasticity of substitution between labor skill groups to $\sigma = 1.3$ based on the review in Card (2009). I set the cost share of commercial floorspace to the estimates from Greenwood, Hercowitz, and Krusell (1997) who measure the share of labor, structures and equipment in value added for the US to be 70, 13, and 17 respectively. A floorspace share of $1 - \alpha_s = 0.156$ corresponds to their estimates renormalized to exclude equipment which is absent from my model. I set this to be equal across industries. I set the elasticity of substitution of demand to $\sigma_D = 6$ close to median estimates from Feenstra et. al. (2014). I vary both elasticities of substitution in robustness checks.

I now discuss how I estimate the parameters $\{\beta, \alpha_{sg}, \kappa, \theta_g, \rho_g, T_g\}$ using relationships from the model.

8.2.2 Parameters Estimated without Solving the Model

Share Parameters I estimate $1 - \beta = 0.24$ to match the long-run housing expenditure share in Bogotá.⁵⁴ I estimate the labor shares α_{sg} by industry using the average share of the wage bill paid to college and non-college educated workers in Colombia between 2000 and 2014 in all cities other than Bogotá. Assuming that firms outside Bogotá aggregate labor using Cobb-Douglas technology, these labor cost shares identify α_{sg} .

Commute Costs The appendix outlines how commute times for car and non-car owners are constructed using averages of the time on each available mode implied by a discrete choice model. In the third-stage of the model, having chosen where to live and work individuals choose which mode of transport to commute with given their idiosyncratic preference for each mode. These preferences are drawn from a Generalized Extreme Value distribution that allows for a nested preference structure across public and private nests. The commute time between each pair of census tracts for each mode is computed in ArcMap. This allows me to estimate κ by fitting the mode choice model to the 2015 Mobility Survey by Maximum Likelihood. This is identified from the sensitivity of individuals' mode choices to differences in times across modes within particular commutes.

Table 7 reports the results. I obtain an estimate of $\kappa = 0.012$, very close to the estimate of 0.01 in Ahlfeldt et. al. (2015). The last entry reports the estimate for the parameter λ governing the correlation of preference shocks within the public nest. When $\lambda = 1$ there is complete independence between draws in the public nest, while $\lambda = 0$ implies perfect correlation. The estimate of $\lambda = 0.14$ suggests idiosyncratic preferences are highly correlated within public nests i.e. commuters on public transport tend to take the quickest mode. Given this order of magnitude difference in sensitivity to commute times, my baseline specifications assume that users take the quickest public mode of transportation available but imperfectly substitute across cars and public transit.⁵⁵ The remaining parameters reflect average preferences for each mode relative to walking (conditional on commute time). Intuitively, cars are most attractive followed by buses and TransMilenio. That TransMilenio is least desirable likely reflects the high crowds using the system as well as the inconvenience of having to walk between stations and final origins and destinations.

⁵⁴See the Engel curves in the appendix.

⁵⁵I explore the sensitivity of my results to alternative aggregation methods in the appendix. The semi-elasticity of mode choices to commute times is κ for public vs private nests and κ/λ for modes within the public nest.

Skill Distribution The gravity equation for commute flows in (2) combined with the specification of commute costs $d_{ija} = \exp(\kappa t_{ija})$ implies a semi-log gravity equation for (conditional) commute flows

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

where γ_{ia} and δ_{jg} are fixed effects and ε_{ija} is an unobserved component of commute times. I aggregate to the locality level and use commuting data from the 2015 Mobility Survey. θ_g is identified from the sensitivity of commuting decisions to commute costs, conditional on trip origin, trip destination and car ownership. Given the presence of zeros in the data, I estimate the model using Poisson Pseudo-Maximum Likelihood (PPML) as suggested by the trade literature (e.g. Santos Silva and Tenrayo 2006).

While the fixed effects absorb any unobserved factors varying by origin and destination, one concern is whether there are origin-destination specific unobserved components of commute costs correlated with trip time. I address this in two ways. First, I include direct measures of factors other than time that may determine the attractiveness of a commute and examine how this changes the results. I measure the average number of crimes along a route, the average house price, as well as the share of the trip that occurs along a main road. Second, I instrument for commute times using the LCP and Tram instruments.

The results are reported in Table 8. Column (1) shows that high-skill workers are less sensitive to commute costs than low-skill workers with a semi-elasticity of -0.0242 compared to -0.0336.⁵⁶ In column (2) I control for other observable factors that may affect the cost of commuting. These characteristics are not significant determinants of commute flows, and the point estimates are unchanged. In columns (3) and (4) I instrument for commute times and find that the estimates are remarkably stable. These results suggest that having controlled for unobservables that vary by origin and destination, commuting decisions are primarily driven by commute times rather than other observed non-time factors. Additionally, the stability of OLS and IV point estimates suggests unobserved determinants of commute costs are not correlated with commute times (conditional on fixed effects).⁵⁷

Given the estimate of κ , the point estimates from column (4) correspond to $\theta_H = 2.054$ and $\theta_L = 2.840$. Both the overall magnitude and the fact that more educated workers are estimated to have a greater dispersion of match-productivities lines up with existing estimates (e.g. Lee 2015; Hsieh et. al. 2016; Galle et. al. 2017).⁵⁸

⁵⁶This does not imply one should observe the high-skilled taking longer commutes. It means that from any location of residence, the low-skilled will be less willing to commute to locations with high commute costs ceteris paribus. Average commute times and distances are greater for the low-skilled in Bogotá since they live further from high employment densities in central areas.

⁵⁷In the appendix I also estimate the relationship using changes in commuting patterns between 1995 and 2015, but there are two reasons why I do not use this as the baseline specification. First, there is a pronounced city-wide reduction in car and bus speeds between 1995 and 2005 discussed in the appendix. While I show this is uncorrelated with TransMilenio routes, a concern is that changes in computed times driven by changes in speeds will be greater for longer trips which introduces an endogeneity problem of its own. Second, IV-PPML failed to converge in the two-period model. Work has shown that the Poisson model is not subject to an incidental parameter problem in the case with two fixed effects (e.g. Fernandez-Val and Weidner 2016), but I am not aware of results for the case with three fixed effects that applies in the time-differenced specification. Regardless, I report the results from the PPML as well as a least squares model (where both OLS and IV estimators converge) in the appendix. The PPML point estimates are very similar to those from the cross-section, and I show in robustness exercises that the quantitative results are qualitatively unchanged when using these alternative θ estimates.

⁵⁸While the only parameter that matters for the computation of equilibria is the ratio $\theta_g = \tilde{\theta}_g / (1 - \rho_g)$, I also able to use wage data to separate $\tilde{\theta}_g$ and ρ_g . I report the results and procedure in the appendix.

8.2.3 Parameters Estimated Solving the Full Model

It remains to estimate the parameters $\{\bar{h}, p_a, T_g, \eta_g, \mu_A, \mu_{U,g}\}$.

In the appendix, I show that given the parameter estimates in the previous section, there is a unique vector of parameters $\{\bar{h}, p_a, T_g\}$ that matches the average expenditure share on housing, the average expenditure on cars, and the college wage premium respectively. I solve for them in the process of recovering the model's unobservables, and allow them to vary in each period to exactly match each wave of data. Lastly, I identify the residential supply elasticity η_g and spillover parameters $\mu_A, \mu_{U,g}$ by exploiting the fact that changes in market access induced by TransMilenio provide a shock to the supply of labor and residents across the city.

Amenities Moment Taking logs of the expression for residential populations in (4) delivers the following expression for residential population growth across skill groups

$$\Delta \ln L_{Riag} = \eta_g \Delta \ln V_{iag} + \eta_g \mu_{U,g} \Delta \ln \frac{L_{RiH}}{L_{Ri}} + \gamma_\ell + \gamma'_R \text{Controls}_i + \Delta \ln \epsilon_{Riag} \quad (18)$$

where $\Delta \ln V_{iag} \equiv \Delta \ln \tilde{y}_{iag} - (1 - \beta) \Delta \ln r_{Ri}$ is the change in indirect utility from living in (i, a) net of changes in amenities, γ_ℓ and Controls_i are locality fixed effects and tract characteristics (to partially control for changing fundamentals) and $\Delta \ln \epsilon_{Riag}$ is a residual that reflects unexplained growth in productivity (i.e. residual variation in $\Delta \ln \bar{u}_{iag}$). To identify η_g , I require a source of exogenous variation in the common component of utility from living in a location $\Delta \ln V_{iag}$. To identify the strength of spillovers $\mu_{U,g}$, I require a separate source of exogenous variation in the college share of residents $\Delta \ln L_{RiH}/L_{Ri}$.

I instrument for the change in indirect utility using the instruments for the change in RCMA. Two additional instruments provide separate variation in the share of college residents. First, tracts which experience a greater growth in CMA to high-skill jobs relative to low-skill jobs should experience a larger increase in the share of college residents. This is captured by the instruments $Z_{Diff,i}^k = \Delta \ln \bar{\Phi}_{RiH}^k - \Delta \ln \bar{\Phi}_{RiL}^k$ for $k \in \{LCP, Tram\}$. Second, I augment this by interacting the change in CMA for high-skilled residents with the house price in the initial period $Z_{Rents,i}^k = \Delta \ln \bar{\Phi}_{RiH}^k \times \ln r_{Ri}^{2000}$. That this should differentially predict entry of high-skilled residents is a direct consequence of log-linearizing the expression for residential populations (4), which implies that poorer, low-skilled residents are less likely to move into expensive neighborhoods due to their increased expenditure on housing.⁵⁹ The moment condition I use to identify η_g and $\mu_{U,g}$ is therefore⁶⁰

$$E[\Delta \ln \epsilon_{Riag} Z_{Riag}] = 0, \quad Z_{Riag} \in \left\{ \begin{array}{ccc} \Delta \ln \Phi_{Riag}^{LCP} & Z_{Diff,i}^{LCP} & Z_{Rents,i}^{LCP} \\ \Delta \ln \Phi_{Riag}^{Tram} & Z_{Diff,i}^{Tram} & Z_{Rents,i}^{Tram} \end{array} \right\}$$

Productivity Moment Recall that firm sales are given by $X_{js} \propto (W_{js}^{\alpha_s} r_{Fj}^{1-\alpha_s})^{1-\sigma_D} A_{js}^{\sigma_D-1}$. Commercial

⁵⁹See footnote 51 for exposition. I define $\bar{\Phi}_{Riag}^k \equiv \sum_a \Phi_{Riag}^k$ to be the sum of RCMA across car ownership within a location-skill group.

⁶⁰I also include orthogonality conditions with each control variable. My baseline specification measures changes in outcomes between 2000 and 2015 and uses the change in transit network due to the first phase of the system since the raw population data at the tract level comes from 2005 (before using the 2015 UPZ totals to inflate to that year).

floorspace prices are observed. Wages are recovered from model inversion in proposition 3 using data on employment, residence and commute costs. These define the labor cost index W_{js} . Lastly, the model implies that firm sales are proportional to the wage bill through $\alpha_s X_{js} = \sum_g w_{jg} \tilde{L}_{Fjgs}$. Since effective labor is obtained using data on employment and model-implied wages, this allows me to recover firm sales X_{js} .

Composite productivity $A_{js} \propto W_{js}^{\alpha_s} r_{Fj}^{1-\alpha_s} X_{js}^{1/(\sigma_D-1)}$ is the residual that ensures the model definition for sales holds. The model infers high productivity in locations where employment is high (reflected through high sales) relative to the observed price of commercial floorspace and the accessibility to workers through the commuting network (which determines wages). Using data before and after TransMilenio's construction provides two values for composite productivities in each location. Taking logs of (12) and including a set of control variables to (partially) capture changing fundamentals yields

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

where $\Delta \ln \epsilon_{Fjs}$ is a residual that reflects unexplained growth in productivity (i.e. residual variation in $\Delta \ln \bar{A}_{js}$).

The agglomeration elasticity is identified from the extent to which model-implied composite productivity depends on employment. The identification challenge is clear: locations may become more productive because more people work there, or locations whose productivity is growing may attract more workers. Guided by the reduced form results, I exploit the fact that labor supply in the model is a log-linear function of FCMA. TransMilenio therefore provides a shock to labor supply in each location through the commuting network, and my instruments isolate the portion of this variation orthogonal to changes in location fundamentals. The moment condition I use to identify μ_A is therefore⁶¹

$$E [\Delta \ln \epsilon_{Fis} Z_{Fig}] = 0, \quad Z_{Fig} \in \left\{ \begin{array}{cc} \Delta \ln \bar{\Phi}_{FiL}^{LCP} & \Delta \ln \bar{\Phi}_{FiH}^{LCP} \\ \Delta \ln \bar{\Phi}_{FiL}^{Tram} & \Delta \ln \bar{\Phi}_{FiH}^{Tram} \end{array} \right\}$$

8.2.4 GMM Results

Main Results Table 9 presents the main results. Three comments are in order. First, the estimate of the productivity externality of 0.237 is large. Ahlfeldt et. al. (2015) obtain an estimate of 0.07 using a similar framework in Berlin, while the estimates in the literature have tended to lie within the 0.03-0.08 range reviewed in the survey by Rosenthal and Strange (2004). However, other experimental approaches in the US have obtained estimates as high as 0.12 and 0.2 (Greenstone, Hornbeck, and Moretti 2010; Kline and Moretti 2014). The returns to agglomeration may be higher in developing countries due to factors such as a lack of road infrastructure or high crime, both of which are certainly at play in Bogotá. To my knowledge, this is the first intra-city estimate of agglomeration in a less developed country using quasi-experimental variation. However, in counterfactuals I

⁶¹Since the unit of observation differs across firm outcomes (tract-industry) and residential outcomes (tract-skill group-car ownership), and there are no interdependencies between parameters across moment conditions (i.e. μ_A only affects the productivity moment condition while $\eta_g, \mu_{U,g}$ only affect the amenity moment conditions), I estimate the equations separately via GMM rather than in one joint system. This avoids the need to make arbitrary aggregation up to consistent units without any loss of information. I also provide a set of 10 different starting values for the solver drawn from a uniform grid between 0 and 10 for each parameter, and find the solver converges to the same minimum each time. This suggests the objective function is well-behaved and my estimates do not result from a local minimum.

turn spillovers off completely as well as set them to a smaller values similar to Ahlfeldt et. al. (2015) to ensure my quantitative results are not driven by this estimate alone.

Second, the residential population elasticity is greater for low-skilled than high-skilled. This might reflect factors such as home ownership that make the residential locations of the high-skilled more sticky. These elasticities are larger than the commute elasticities, underscoring the benefits from estimating residential and employment location decisions as a two stage problem.

Third, the spillover parameters for residential amenities are 0.250 and 0.342 for the low- and high-skilled respectively. The share of college-educated residents in a tract increases the amenities from living there, the high-skilled value living around each other more than the low-skilled, and these endogenous forces are around twice as strong in Bogotá than Ahlfeldt et. al. (2015) find for Berlin.

Robustness In the appendix, I check the robustness of these estimates. First, I control for log distance to the closest TransMilenio station (instrumented using the log distance to the instruments) to assess whether TransMilenio impacted amenities and productivities directly (e.g. through street improvements, effects on crime or pollution). I find no effect on amenities,⁶² and an economically small reduction in productivity. This suggests the primary effect of TransMilenio was indeed through improved accessibility. Next, I vary the elasticity of substitution of demand (from 4 to 9) and elasticity of substitution between skill groups (from 1.3 to 2.5).⁶³ The point estimates are largely robust to the elasticity of substitution of labor, but the agglomeration point estimate is mechanically related to the demand elasticity since both affect returns to scale. While my preferred estimate lies in the middle of the observed range, I later examine the robustness of my quantitative results to these parameters. Finally, I assess how excluding non-homothetic housing demand impacts my estimates which I discuss in Section 9.

8.3 Non-targeted Moments: Model vs Data

I now evaluate the model's performance by assessing its ability to match moments not targeted in estimation.

Wages Figure 5a compares the average wage earned by residents of each locality with that observed in 2014 in the GEIH data. The latter was not used in estimation. The two variables are highly correlated with values of 0.528 for non-college and 0.592 for college workers. Most observations lie along the 45-degree line for low-skilled workers, but there is noticeable deviation for the richest localities amongst high-skill workers.⁶⁴ While the model is unable to capture all factors that drive differences in average income, the high correlation suggests that the spatial forces perform well in explaining income differences across the city.

⁶²In additional results available on request, I show directly there was no significant impact of distance to TransMilenio on change in crime between 2007 and 2012.

⁶³The value of 2.5 is estimated by Card (2009) for skill groups using regional data in the US.

⁶⁴The model loads all spatial variation in employment unexplained by bilateral commute costs from places of residence into wages. There may be non-wage amenities such as safety, access to food and retail establishments, that determine how attractive it is to work in a location. This could explain the relatively poor fit for high-skilled workers if they value these attributes more. So long as these amenities are fixed in the counterfactuals they will not affect the quantitative results.

Amenities and Productivities In the model, amenities and productivities represent characteristics that make locations more or less desirable to individuals and firms who might choose to locate there. In the appendix, I show that (i) neighborhoods with less crime are associated with higher amenities and (ii) productivities are higher in tracts with less crime, a flatter slope and a higher density of roads. Overall, the model performs well at capturing features that affect the desirability of locations in the city.

Commute Flows Using the wages that rationalize the observed distribution of residential population by skill and employment by industry, I use the model to compute the commute flows between origin, destination and car ownership pairs according to the gravity equation (2). The model exactly matches total residence, employment and the share of car owners, but information over the bilateral commute shares was not used in estimation. Figure 5b tests the performance of the model's assumptions by comparing these implied commute shares with those observed within each locality in the 2015 Mobility Survey. The model matches commute flows across skill and car ownership groups well. Importantly, the fit is even across college and non-college workers, showing the method to back out wages by skill group using the location of employment by industries performs well in predicting commute flows.

Employment By Skill Group Figure 6 compares the skill employment ratio $\ln(L_{FiH}/L_{FiL})$ within each UPZ in the model with that implied by trips to work in the 2015 Mobility Survey. To show the importance of the ingredients in the model, panel (a) plots the results from a simplified model in which labor skill groups are perfect substitutes and share the same commute elasticity (set to the average value) as in Ahlfeldt et. al. (2015). In this model, relative employment by skill group has an oddly smooth pattern that slowly declines as one moves further south in the city. This is because low- and high-skilled workers receive the same relative wage in each location and have the same sensitivity to commute costs, so differences in commuting behavior are solely due to differences in residential locations. Thus, the supply of high-skilled workers is much greater in the North close to where they live, and vice versa for the poor who live in the South. This is clearly counterfactual to the distribution in the data shown in panel (c). By contrast, my model performs much better in matching this spatial distribution of the employment of relative skills (panel (b)): the correlation between the skill share in the data and in the baseline model is 0.406 compared to 0.256 in the simplified model.

Heterogeneous Sorting Response Panel (c) of Figure 7 shows that when simulating the change in transit induced by TransMilenio, the model captures the increase in residential segregation documented in the reduced form analysis: the college share rises (falls) with RCMA in tracts with a high (low) initial college share. Running the same regression as the last column of Table 5 yields an interaction coefficient between the change in average RCMA and the initial college share measure of 0.102 (0.020), remarkably close to the 0.111 (0.047) estimated in the data despite not being targeted in estimation.

9 The Welfare Effects of TransMilenio

This section quantifies the impact of TransMilenio by simulating the effect of its removal from the 2012 equilibrium. I begin by using first-order approximations to the welfare change that are easy to implement using readily available travel survey data. I then compare this with the results from the full model that capture the global, general equilibrium responses and decompose the differences.⁶⁵

9.1 First Order Effects

The standard approach to evaluate the gains from transit infrastructure is based on the Value of Travel Time Savings (VTTS) approach (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.

Proposition 4. *In the model without spillovers, equal ownership of the housing stock and no preference heterogeneity, the competitive equilibrium is efficient. The elasticity of welfare to a change in commute costs is*

$$d \ln \bar{U}_g = -\beta \kappa \sum_{i,j,m} \frac{w_{ijg} L_{ijgm}}{\sum_{r,s,n} w_{rsg} L_{rsgn}} dt_{ijm}. \quad (19)$$

A straightforward application of the envelope theorem shows that, in an efficient economy, only the direct effect of new infrastructure matters for welfare. This turns out to be proportional to a weighted average of commute time reductions, with weights given by a commute's initial share of total labor income. The constant of proportionality reflects the relation between commute time and costs (through κ), and that with fixed housing supply a reduction in commute times raises house prices and hence the cost of living (at rate $1 - \beta$).⁶⁶

In the appendix, I derive two additional welfare approximations:

$$d \ln \bar{U}_g^{VTTS} = -\kappa \sum_{i,j,m} \frac{L_{ijgm}}{\sum_{r,s,n} L_{rsgn}} dt_{ijm}$$

$$d \ln \bar{U}_g^{FOA} = \sum_{i,m} \frac{L_{Rigm}}{\sum_{r,n} L_{Rr gn}} (\varepsilon^Y - \beta_{img} \varepsilon^R) d \ln \Phi_{Rigm}$$

The first is the welfare change under the VTTS approach. This mirrors (19), except for the slight difference in weights and the omission of the house price response captured by β . The second takes a less model-dependent approach that leverages the reduced form elasticities estimated above. The expression arises from a first order approximation to the change in an indirect utility function that depends only on house prices and income,

⁶⁵I refer to the “2012 equilibrium” as the post-TransMilenio equilibrium. Population and employment data come from 2015, land market data come from 2012, and the TransMilenio network is taken to include phases 1 and 2 of the system. Underlying data for first order approximations comes from 2015 mobility survey. See the appendix for additional information on the algorithm used to solve for counterfactual equilibria. As highlighted by Proposition 1, there may be multiple equilibria in the presence of spillovers. The selection rule I use 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.

⁶⁶This occurs despite the fact that individuals own their homes. Suppose commute costs fall by 1%. RCMA and house prices increase by 1%, wages are unchanged, so income increases by 1% while the cost of living rises by $(1 - \beta)\%$, implying a $\beta\%$ rise in welfare.

where $d \ln \Phi_{Rigm}$ is a simple approximation to the change in RCMA using initial employment levels and the change in travel times.⁶⁷ The elasticities $\varepsilon^X \equiv \frac{\partial \ln X}{\partial \ln \Phi_{Ri}}$ are the reduced form IV estimates above and β_{img} is the expenditure share on housing, approximated using the group average across the city that I observe in the data. Ignoring the indirect effects of changing commute costs on employment and FCMA yields $d \ln \Phi_{Rigm} \approx -\theta \kappa \sum_j \pi_{j|img} dt_{ijm}$. Thus, this approximation is also proportional the average time savings with constant of proportionality $\theta \kappa (\varepsilon^Y - \beta_g \varepsilon^R)$.

Results Table 10 presents the results. Each entry reports the absolute value of the welfare impact of removing TransMilenio from the post-period equilibrium. I begin by comparing the first order approximations in rows one to three, which differ mainly due to the different constants of proportionality used.

In row one, the VTTS welfare gains are around one third higher than in the model's first order approximation since $\beta = 0.76$. The model's approximation captures that reductions in commute times raise housing prices and hence the cost of living.

In row two, using the CMA elasticities the benefits of TransMilenio are only 40% of those using the model's approximation since $\theta (\varepsilon^Y - \beta_g \varepsilon^R) \kappa \approx 0.402 \times \beta \kappa$. Since the estimate of $\varepsilon^R = 0.557$ is very close to its implied value 0.5608 in the model used to generate (19), this difference is driven primarily by the estimate for ε^Y . The model implies $\varepsilon^Y = 1/\theta \approx 0.384$, while the estimate from the data is only 0.282. This is because in reality, only a portion of the gains will manifest through more time for work (and higher incomes) while the remainder will accrue through increased leisure time. Unfortunately only the elasticity of the former to CMA is easily measured in the data. By contrast, in the model all saved time is used for labor and thus translate fully into higher incomes. An alternative model (considered below) in which commute costs reduce utility rather than effective labor would have the same elasticity $1/\theta$ for the direct effect of higher CMA on welfare even though there would be no direct effect of saved time on incomes. One can therefore think of $1/\theta$ as the model's way of capturing the full direct welfare benefits of improved CMA on utility.⁶⁸ Note than in all cases, the low-skilled benefit more reflecting their greater reliance on public transit.

How do the welfare effects implied by these approximations differ with those that capture the full general equilibrium response? Comparing rows 3 and 4 shows that the level of welfare gains is smaller in the full general equilibrium model while the effect on inequality is reversed. There are a number of factors that could be driving these differences. First, the approximation holds all discrete choices fixed while the GE models incorporates the substitution elasticities that determine the ability of agents to reorganize spatially. This may be important for large changes in commute costs. Second, the equilibrium responses may now be important due to the presence of externalities and the large size of the shock. Third, there is approximation error in measurement of the direct effect.⁶⁹ I now turn to decomposing the sources of these differences.

⁶⁷Specifically, I approximate $d \ln \Phi_{Rigm} \approx \ln \left(\sum_j (d'_{ijm})^{-\theta} L_{Fjg} \right) - \ln \left(\sum_j (d_{ijm})^{-\theta} L_{Fjg} \right)$ which ignores the FCMA term in the denominator and holds employment fixed at its initial level. A similar first order approach is used in Atkin et. al. (2018).

⁶⁸Indeed, when replacing $\varepsilon^Y = 1/\theta$ the benefits from TransMilenio implied by the CMA elasticities become 75% of those in the model-based approximation. See the appendix for the precise cluster of structural parameters that make up ε^R ; I assume no amenity spillovers when making the comparison since this is what is assumed in the model-free approximation.

⁶⁹The gross percentage change in commute costs is $\hat{d}_{ij} = \exp(-\kappa dt_{ij})$ where $dt_{ij} > 0$ is the increase in commute times due to removing TransMilenio. Thus $d \ln d_{ij} = -\kappa dt_{ij}$ used in the first-order approximations is an overestimate of the true effect. While the

9.2 Unpacking the General Equilibrium Response

Aggregate Effects Table 11 presents the effect of TransMilenio on GDP, total rents and welfare. Panel A presents the closed city results, in which the population of the city remains constant and utility adjusts in equilibrium. The effects on all outcomes are large, independent of whether spillovers are included: TransMilenio increases city GDP between 3.12%-3.92%, total city rents by 3.29%-3.72% and worker welfare by around 3.5-3.9%, the higher number referring to the case with spillovers. Panel B shows the open city results, where population rather than utility adjusts (which is fixed to the reservation level in the wider economy). The effects of TransMilenio are substantial, increasing the population of the low-skilled by 8.56%-10.74% and the high-skilled by 9.54%-12.30%. The increase in factor supply leads to greater GDP growth between 10.34%-15.59%. The population influx fuels greater house price appreciation, with total rents rise between 13.15-16.28% due to TransMilenio.

TransMilenio improved the spatial allocation of employment and residence. Figure 7 plots the change in employment and population in each tract by each variable's initial level. Panel (a) shows that tracts with the largest employment lose the most when TransMilenio is removed. By enabling productive locations with high employment to grow the most, the system's efficiency gains are driven in part by an improvement in the spatial allocation of labor. Panel (b) shows similar patterns hold for residence, but the effects are more muted.

Costs vs Benefits How did the output gains from TransMilenio compare with the costs of the system? Panel A of Table 13 provides a breakdown of the costs and benefits of the system (see appendix for details on cost calculations). Even using the most conservative estimate in column (1), I find that the net present value of the net increase on GDP was about \$50bn, or a net increase of 2.73% in the steady-state level of GDP. This suggests the system was a highly profitable investment for the city.

Distributional Effects The previous results document that welfare inequality (defined as \bar{U}_H/\bar{U}_L) rises in the general equilibrium model, the opposite result to the first order welfare impact. Table 12 decomposes the channels through TransMilenio affects welfare in the full model, starting with a simplified case of the model and slowly adding in the model's ingredients to unpack this final result.⁷⁰

Rows 1 to 4 consider the model where workers share the same (average) value for η and θ and are perfect substitutes in production. Relative wages are therefore equal employment locations. In the first row, I simulate the effect of removing TransMilenio, and then adding it back in holding all discrete choices and prices constant. Individuals gain only through time saved on their commute. Low-skilled workers benefit the most with inequality falling by 0.89%, both because they rely more on public transit and live in peripheral locations where

sign of the approximation error flips if commute times fall rather than rise, its magnitude is the same (for the range of time changes implied by TransMilenio). There are also potential differences due to data aggregation, since the travel surveys are aggregated to the UPZ-level while the GE model is at the census-tract level. Lastly, recall the general equilibrium results in Table 10 are from the model where each resident owns an equal share of housing stock resided in by their type, and thus differs from the baseline results where Table 11 where all workers are renters.

⁷⁰For each model, both the baseline and counterfactual equilibria are re-computed. This involves solving for updated values for T_H, \bar{h}, p_a and w_{jg} . The results from the decomposition differ slightly when spillovers are set to their estimated values, since the counterfactual equilibrium without TransMilenio is different.

access to jobs improves the most.

The second row allows individuals to change their choices of where to live, work, and whether to own a car, but keeps prices fixed. Allowing for more margins of adjustment increases the level of welfare gains, which accrue disproportionately to the low-skilled with inequality falling by 1.35%. In the third row, I allow for general equilibrium adjustment of house prices. This reduces the level of welfare gains which now are shifted from workers to landlords. House price appreciation hurts the low-skilled the most, precisely because prices rise the most in outlying areas where the accessibility gains are greatest. The fourth row allows for full general equilibrium adjustment yielding qualitatively similar results. While the effects on inequality are the same in partial and general equilibrium, the difference in the level of welfare gains shows that equilibrium effects start to matter with large changes in the presence of spillovers.

In the fifth row, I allow workers to differ in their commute elasticities set to the estimated values. In the first order approximation results, discrete choices over where to work are fixed. In the full model, individuals can substitute between commutes in response to the removal of TransMilenio with elasticity of substitution θ_g . Since low-skilled workers have a greater commuting elasticity than the high-skilled, they are able to substitute to less costly commutes more than high-skilled workers whose commuting choices stickier. Intuitively, the group with the lower elasticity bears a greater incidence of slow transit infrastructure through this channel, and this shifts the gains towards the high-skilled with welfare inequality now falling by only 0.19%. The sixth row incorporates differences in residential choice elasticities, and the result is qualitatively similar.

The final row considers the full model with imperfect substitution in production. Relative wages now differ for each group across employment locations, based on demand and supply in each tract's labor market. This has two effects. First, what matters is whether workers are connected to locations where demand for their specific skill (and hence their wage) is highest. For the geography of Bogotá and TransMilenio, this tends to benefit the high-skilled who are live in concentrated part of the city in the North which TransMilenio connected with the high skill-intensive industries in the center and center-north (Figures 2 and 3). Residence and employment for the low-skilled is more dispersed, so TransMilenio connected a smaller fraction of these workers with high-wage locations. Second, since skill groups are imperfect substitutes in production, high-skilled workers are now partially shielded from the reduction in wages due to the large labor supply shift of low-skilled workers who use public transit the most. These forces mean that ultimately TransMilenio increased welfare inequality between the low- and high-skilled increased by 0.369%.

Taken together, I find that the incidence of improving public transit depends not only on how much each group uses it, but also how willing each group is to bear high commute costs to work at a particular location, whether the system connects workers with high-wage locations and the general equilibrium response of wages and house prices. While the first order approximation slightly overestimates the aggregate welfare effect, it delivers the opposite effect on equity by only considering the first channel. My results therefore imply that while there is no universal answer to who benefits most from transit infrastructure (which depends on the geography of the city and the transit improvements), accounting for spatial reorganization of the city and general equilibrium adjustment of prices implies that investments in public transit are a less precise way to target welfare improvement for the poor than is implied by the travel time savings approaches.

9.3 Robustness and Model Extensions

In the appendix, I explore the robustness of my quantitative results to alternative parameter values. The effects on output, rents and welfare are qualitatively similar across specifications.⁷¹

In Table 14, I explore the sensitivity of the quantitative results to a number of model extensions. First, I assume that the shocks by workplace location affect preferences rather than productivity.⁷² In this model, TransMilenio no longer acts as a positive supply shock to each location (holding employment decisions constant). As a result, wages do not fall via this channel and welfare gains from the system rise. Since labor supplied by each worker is unchanged by commute costs, the effect on output falls by more than two thirds. However, this difference is eliminated by the increase in labor supply from population growth in the open city model.

Second, I allow workers to make a joint decision over where to live and work.⁷³ The results are qualitatively unchanged from the baseline model, suggesting the timing assumption has little quantitative impact.

Third, I confront the fact that neither census nor CCB employment data cover employment in domestic services. From 2000-2014, 7.3% of non-college educated Bogotanos worked as domestic helpers while almost no college educated workers did. On the one hand, the model may under-estimate the welfare gains to the low-skilled by ignoring the fact that TransMilenio likely improved access to domestic services jobs in the homes of the college educated in the North. On the other hand, high-skilled workers also benefit from this increased labor supply which lowers the cost of hiring domestic services. In the appendix, I extend the baseline model to include employment in domestic services and outline its calibration. The fifth row of Table 14 shows that allowing for employment in domestic services has little effect on the distributional effects: the benefits for low- and high-skilled workers roughly balance out.

Fourth, I incorporate home ownership. By assuming all individuals are renters, I may understate welfare gains to homeownership low-skilled individuals who disproportionately benefit from appreciation in the city's outskirts. At the same time, the high-skilled are more likely to own their home with an ownership rate of 0.603 compared to 0.457 for the low-skilled in 2015. In the appendix, I extend the model as follows. I assume local land rents are collected into a fund, a fraction of which is re-distributed to low- and high-skilled workers based on the home ownership rates. The remainder is paid into an aggregate portfolio which is owned in equal shares by all residents. The final row of Table 14 shows that home ownership does indeed close the gap in welfare gains between high- and low-skilled workers, but not enough to reverse it.⁷⁴

⁷¹These include increasing the values of θ and η by 50% (in case my estimates reflect medium- rather than long-run responses), using alternative values of θ estimated via PPML in two periods, setting spillovers to one third of their estimated values (to match the magnitude of productivity spillovers in Ahlfeldt et. al. 2015), using a larger elasticity of substitution across labor skill groups $\sigma_L = 2.5$, measuring the distribution of employment using the 2005 census rather than the 2015 CCB (to address whether missing informal establishments impacts the results), using alternative values of the elasticity of demand $\sigma = 3, 9$, and allowing for preference heterogeneity within public transit modes by setting $\lambda = 0.14$ from Table 7.

⁷²In the preference shock and joint decision model, I assume that there are no fixed costs i.e. $p_a = \bar{h} = 0$. Rows 3 and 4 should therefore be compared with row 2.

⁷³I assume that workers draw a joint preference for each residence-employment pair drawn from a Fréchet with shape θ_g . I estimate θ_g from the implied gravity equation for unconditional commute flows, finding that $\hat{\theta}_L = 3.058$ and $\hat{\theta}_H = 1.772$.

⁷⁴The effect of TransMilenio on output is greater in this model since house price appreciation increases expenditure by residents who own the housing stock, instead of being spent outside of the city by absentee landlords. The effect on welfare is attenuated since income now depends on the sum of labor income and income from home ownership, and the direct effect of TransMilenio is greater on the former in proportionate terms.

Role of Non-homotheticities While a homothetic model can match the same data by loading differences in observed choices into exogenous preference shifters, my results suggest three reasons why allowing for non-homotheticities is important. First, the first two rows of Table 14 imply that while the two models yield similar results for the average welfare effect, the homothetic model underestimates the increase in inequality by 35%. This is due to the low-skilled being hurt more by house price appreciation when they spend a greater share of their income on it. Second, panel (c) in Figure 7 shows that non-homotheticities are necessary to match the high-skilled sorting into relatively higher amenity neighborhoods in response to an improvement in accessibility. Assessing this quantitatively, in the data the interaction between the change in average RCMA and the initial college share is 0.111 (0.047) in Table 5, compared with 0.102 (0.020) in the model with non-homotheticities but only 0.06 (0.017) without them. This differential response to improved accessibility across neighborhoods is precisely the rationale behind the instruments in Section 8. Third, in the estimation robustness exercise reported in the appendix, I show that ignoring non-homotheticities in housing expenditure shrinks the estimates of the amenity spillover by about one quarter with implications for efficiency.⁷⁵

10 Policy Counterfactuals

10.1 Impact of Different Lines and Planned Cable Car

Table 15 evaluates the effects of different portions of the system. Row (1) evaluates the impact of adding a Cable Car to the slums in the hills of Ciudad Bolivar in the South. The aggregate effects of the line are small due to its modest size, but it benefits the low-skilled workers who are more likely to live in targeted areas. Rows (2) and (3) simulate the effect of removing lines H and A that connect the South and North of the city with the CBD. Both lines drive large improvements in welfare, but these accrue disproportionately to the group that lives along it (the low- and high-skilled, respectively). This underlines that who benefits from new commuting infrastructure depends on the geography of the city and transit network in question. Row (4) shows that the feeder system, which connects outlying areas with portals using buses that run on existing roadways, increases welfare more than any other line of the network. This underscores the large benefits to providing cheap, complementary services that reach residents in outlying but dense residential areas, thereby reducing the last-mile problem of traveling between stations and final destinations. Row (5) isolates the contribution of route placement to TransMilenio’s welfare gains by applying a uniform reduction in times across all commutes and modes equal to the average percentage change from the system.

10.2 Land Value Capture

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 where it was needed. I show in the appendix that housing supply did not respond to the system’s construction, consistent with other evidence on the restric-

⁷⁵When $\bar{h} > 0$, one can use the amenities moment condition (18) to show that if the non-homotheticity is ignored in estimation, the estimate of $\eta_g \mu_{U,g}$ will be downward biased when the growth in the college share is positively correlated with the growth in house prices, which is true in the data.

tive 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 construction of the system.

I evaluate the impact of TransMilenio if housing supply had responded to the opening of the system. In the most extreme case, I assume that housing supplies freely adjust to reflect long-run adjustment. This provides a useful upper bound on welfare gains from facilitating housing supply response. I then simulate the effect of two potential LVC schemes implemented by the government. 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.⁷⁶ 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.⁷⁷ Details on these model extensions are provided in the appendix. I compare the two equilibria from first removing TransMilenio and then adding it back under each housing supply model.

Table 16 presents the results. In the closed city model, welfare increases by 24.4% more under free adjustment. Under the LVC schemes, welfare improves by 23% and 6.9% under the CMA and distance-band policies respectively. Taking the CMA scheme as an example, welfare rises by 24.9% for low-skilled individuals versus 18.3% for the high-skilled. More accommodative zoning benefits the low-skilled by dampening house price appreciation towards the edge of the city where they live. While this exercise does not factor in potential adjustment costs associated with demolitions and displacement of prior residents, it highlights the large potential benefits left on the table. In terms of government revenues, the distance-band based permit measure recoups only 9.9% of the capital cost of the system, compared to the 17.9% earned using the commuter market access-based permit in the closed city case. In the open city, these increase to 27% and 50% respectively.⁷⁸

My 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. Additionally, the comparison with the distance-based policy underscores how measures of CMA can be used as parsimonious tools for governments to guide the allocation of rezoning. By targeting central and outlying areas where accessibility for residents and firms increase the most, the CMA-based policy concentrates new construction where it is most needed.

⁷⁶These schemes have a number of benefits over property taxes. They are 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. While the precise increase is hard to find, revenues from the scheme covered 20.5% of costs. In examples covered in Salon (2014) between 14-88% of costs are covered. The 30% increase in permitted densities I choose therefore results in similar revenues. Of course, the revenue raised varies across alternative candidate policies.

⁷⁷In particular, I let the change in permitted FAR be proportional to $\vartheta_i \Delta \ln \Phi_{Ri} + (1 - \vartheta_i) \Delta \ln \Phi_{Fi}$ where ϑ_i are the residential floorspace shares in the initial equilibrium and $\Delta \ln \Phi$ are the instruments for the change in CMA holding population and employment at their initial values. Each of these values is based only on information the government has at the time of the policy change.

⁷⁸Panel B in Table 13 shows the revenues are qualitatively similar with spillovers shut down.

11 Conclusion

This paper makes three contributions to our understanding of the aggregate and distributional effects of urban transit systems. First, I build a quantitative general equilibrium model of a city where low- and high-skill workers with non-homothetic preferences sort over where to live, where to work, and whether or not to own a car. Second, I show this theory provides a reduced form framework to evaluate the effects of transit based on “commuter market access”. Third, I use the construction of the world’s largest BRT system in Bogotá, Colombia to show this framework parsimoniously explains the change in city structure, and quantify the effects of the system and counterfactual policies.

I find the reduced form representation fits the heterogeneous adjustment of population, employment and housing markets to TransMilenio. While the system caused increases in welfare and output larger than its cost, the gains accrued slightly more to high-skilled workers. The opposite distributional impact is predicted by the standard value of time savings approach that ignores spatial reorganization of the city and general equilibrium adjustment of prices that occur from large changes in infrastructure. My findings therefore imply that investments in public transit are a less precise way to target welfare improvements for the poor than is implied by the typical framework. I also find the welfare gains would have been around one fourth larger had the government implemented a more accommodative zoning policy, underscoring 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. (2011), “If Alonso was Right: Modeling Accessibility and Explaining the Residential Land Gradient”, *Journal of Regional Science*, 51(2): 318-338.
- 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.
- AHLFELDT, G.M., REDDING, S.J., AND D.M STURM (2016), “A Quantitative Framework for Evaluating the Impact of Urban Transport Improvements,” Working Paper.
- 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 (2017), “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., AND C. ARKOLAKIS (2019), “The Welfare Effects of Transportation Infrastructure Improvements,” Working Paper.
- ALLEN, T., ARKOLAKIS, C. AND X. LI (2015), “Optimal City Structure,” Working Paper.
- ALLEN, T., ARKOLAKIS, C. AND Y. TAKAHASHI (2017), “Universal Gravity,” Working Paper.
- ALONSO, W. (1964), *Location and Land Use*. Harvard, Cambridge MA.
- ATKIN, D., FABER, B., AND GONZALEZ-NAVARRO, M. (2018), “Retail globalization and household welfare: Evidence from Mexico”, *Journal of Political Economy*, 161(1): 1-73.
- ANAS, A., ARNOTT, R., AND K. SMALL (1998), “Urban Spatial Structure”, *Journal of Economic Literature*, 36(3): 1426–64.
- BARTELME, D. (2015), “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., 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.
- BOCAJERO, 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.
- BRT DATA (2017), *Global BRT Statistics*, Retrieved from <http://brtdata.org>.
- BRYAN, G., AND M. MORTEN (2018): “The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia,” *Journal of Political Economy*, forthcoming.
- CARD, D. (2009). “Immigration and Inequality” *American Economic Review*, 99(2): 1–21.
- CERVERO, R. (2013), “Bus Rapid Transit (BRT): An Efficient and Competitive Mode of Public Transport”, *Berkeley Institute of Urban and Regional Development Working Paper 2013-01*.
- 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.
- CICCONI, A., AND R. HALL (1996), “Productivity and the Density of Economic Activity,” *American Economic*

Review, 86(1), 54–70.

- 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.
- COUTURE, V. (2016), “Valuing the Consumption Benefits of Urban Density”, Working Paper.
- CRACKNELL, J. (2003), “Transmilenio Busway-Based Mass Transit, Bogotá, Colombia”, World Bank Working Paper.
- DAVIS, L. (2008), “The Effect of Driving Restrictions on Air Quality in Mexico City”, *Journal of Political Economy*, 116(1):38-81.
- 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. (Forthcoming) “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, forthcoming.
- DONALDSON, D. (2015), “The Gains from Market Integration”, *Annual Review of Economics*, 7:619-647.
- DONALDSON, D., AND R. HORNBECK (2015), “Railroads and American Economic Growth: A Market Access Approach”, *Quarterly Journal of Economics*, 131(2): 799-858.
- DURANTON, G., AND M. TURNER (2012), “Urban Growth and Transportation”, *Review of Economic Studies*, 79: 1407-1440.
- DURANTON, G., AND M. TURNER (2017), “Urban form and driving: Evidence from US cities”, Working Paper.
- EATON, J., AND S. KORTUM (2002), “Technology, Geography, and Trade,” *Econometrica*, 70(5), 1741–1779.
- FAJGELBAUM, P. D., AND E. SCHAAL (2017), “Optimal Transport Networks in Spatial Equilibrium,” NBER Working Paper, 23200.
- FEENSTRA, R., P. LUCK, M. OBSTFELD AND K. RUSS (2014), “In Search of the Armington Elasticity,” NBER Working Paper No. 20063.
- FERNANDEZ-VAL, I. AND M. WEIDNER (2016), “Individual and Time Effects in Nonlinear Panel Data Models with Large N, T,” *Journal of Econometrics*, 196: 291-312
- FORTHERINGHAM, S., AND M. O’KELLY (1989), *Spatial Interaction Models: Formulations and Applications*, Kluwer, Dordrecht.
- GADUH, A., GRACNER, T., AND A ROTHENBERG (2018) “Improving Mobility in Developing Country Cities: Evaluating Bus Rapid Transit and Other Policies in Jakarta”, Working Paper.
- GALLE, S., RODRIGUEZ-CLARE, A., AND M. YI (2017), “Slicing the Pie: Quantifying the Aggregate and Distributional Consequences of Trade”, Working Paper.
- GIANNONE, E. (2018), “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 (2016), “Subways and urban growth: Evidence from Earth”, Working Paper.
- GUERRIERI, G., HARTLEY, D., AND E. HURST (2013), “Endogenous Gentrification and Housing Price Dynamics” *Journal of Public Economics*, 100:45-60.
- GRAFTIEAUX, P., AND D. HIDALGO (2005), “A Critical Look at Major Bus Improvements in Latin America and Asia: Case Study TransMilenio, Bogotá, Colombia”, Working Paper.

- 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.
- GREENWOOD, J., HERCOWITZ, Z., AND P. KRUSELL (1997): “Long-Run Implications of Investment-Specific Technological Change,” *American Economic Review*, 87(3), 342-362.
- HANDY, S., AND D. NIEMEIER (1997), “Measuring Accessibility: An Exploration of Issues and Alternatives” *Environment and Planning A*, 29(7): 1175— 94.
- HANSEN, W. (1959), “How Accessibility Shapes Land Use,” *Journal of the American Institute of Planners*, 25: 73–76.
- HARRIS, C. (1954), “The Market as a Factor in the Localization of Industry in the United States,” *Annals of the Association of American Geographers*, 44(4), 315–348.
- HEILMANN, K. (2016), “Transit Access and Neighborhood Segregation:A Study of the Dallas Light Rail System”, Working Paper.
- HONG, Y-H., MURAKAMI, J., SUZUKI, H., AND B. TAMAYOSE (2015), *Financing Transit-Oriented Development with Land Values*, Washinton, DC: The World Bank.
- HEBLICH, S., REDDING, S., AND D. STURM (2018): “The Making of the Modern Metropolis: Evidence from London,” Working Paper.
- HSIEH, C-T., HURST, E., JONES, C., AND P. KLENOW (2016), “The Allocation of Talent and US Economic Growth”, Working Paper.
- 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., (2016), “The Effects of License Plate-Based Driving Restrictions on Air Quality: Theory and Empirical Evidence”, Working Paper.
- LEE, E. (2015), “Trade, Inequality, and the Endogenous Sorting of Heterogeneous Workers”, Working Paper.
- LLERAS, G. (2003), “Bus Rapid Transit: Impacts on Travel Behavior in Bogota”, MIT Masters Thesis.
- LUCAS , R., AND E. ROSSI-HANSBERG (2002), “On the Internal Structure of Cities,” *Econometrica*, 70(4), 1445–1476.
- MACKIE, P., J. NELLTHORP, AND J. LAIRD (2005), “Notes on the Economic Evaluation of Transport Projects”, World Bank Transport Note No. TRN-15.
- MAJID, H., A. MALIK, AND K. VYBORNÝ (2018): “Infrastructure Investments, Public Transport Use and Sustainability: Evidence from Lahore, Pakistan,” Working Paper.
- MCARTHUR, D., L. OSLAND, AND I. THORSEN (2012), “Spatial Transferability of Hedonic House Price Functions”, *Regional Studies*, 46(5): 597-610.
- MCDONALD, J., AND D. MCMILLEN (2010), *Urban Economics and Real Estate: Theory and Policy*, John Wiley & Sons, Hoboken, NJ.
- 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.
- MILLS, E. (1967), “An Aggregative Model of Resource Allocation in a Metropolitan Centre,” *American Economic Review*, 57(2), 197–210.
- MONTE, F., REDDING, S., AND E. ROSSI-HANSBERG (2017), “Commuting, Migration and Local Employment Elasticities”, NBER Working Paper 21706.
- MUTH, R. (1969), *Cities and Housing*, Chicago: University of Chicago Press.
- OWENS, R., E. ROSSI-HANSBERG, AND P-D. SARTE (2017), “Rethinking Detroit”, Working Paper.
- REDDING, S. (2010), “The Empirics of New Economic Geography”, *Journal of Regional Science*, 50(1), 297-311.
- REDDING, S., AND A. VENABLES (2004), “Economic Geography and International Inequality,” *Journal of Inter-*

national Economics, 62, 53–82.

- REDDING, S., AND D. STURM (2008), “The Costs of Remoteness: Evidence from German Division and Reunification,” *American Economic Review*, 98(5), 1766–1797.
- REDDING, S., AND D. STURM (2016), “Estimating Neighborhood Effects: Evidence from War-time Destruction in London”, Working Paper.
- 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.
- RUIZ, F., AND G. VALLEJO (2015), “Using Land Registration as a Tool to Generate Municipal Revenue: Lessons from Bogota”, Cadastre of Bogota Working Paper.
- SALON, D. (2014), “Location Value Capture Opportunities for Urban Public Transport Finance,” Transit Leadership Summit, London, White Paper.
- SANTOS SILVA, J.M, AND S. TENRAYO (2006), “The Log of Gravity,” *The Review of Economics and Statistics*, 88(4), 641–658.
- SEVEREN, C. (2016), “Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification”, Working Paper.
- SMALL, K., AND E. VERHOEF (2007), *The Economics of Urban Transportation*, New York: Routledge.
- URIBE SANCHEZ, M.C. (2015), “Land Information Updating, a De Facto Tax Reform: Bringing up to Date the Cadastral Database of Bogota”, Cadastre of Bogota Working Paper.
- WRIGHT , L., AND W. HOOK (2007), *Bus Rapid Transit Planning Guide*, New York: Institute for Transportation and Development Policy.

Tables

Table 1: College-Employment Shares by Industry

Industry	College Share	Employment Share
Domestic Services	0.085	0.050
Construction	0.181	0.052
Hotels & Restaurants	0.235	0.057
Wholesale, Retail, Repair	0.300	0.222
Manufacturing	0.315	0.173
Transport, Storage, Communications	0.341	0.089
Other Community, Social, Personal Serv	0.380	0.050
Real Estate	0.556	0.120
Social & Health Services	0.634	0.053
Public Administration	0.707	0.038
Education	0.810	0.052
Financial Services	0.827	0.028

Note: Data is an average over 2000-2014 and comes from the GEIH and ECH. The first column shows the share of workers which have post-secondary education within each one-digit industry. The second column shows the industry's share of total city employment. Only industries accounting for at least 1% of employment reported.

Table 2: Commuting in 1995

	lnSpeed	lnSpeed	Bus	Bus
Bus	-0.353*** (0.021)	-0.305*** (0.016)		
Low-Skill			0.287*** (0.010)	0.163*** (0.011)
R^2	0.06	0.76	0.18	0.47
N	14,841	12,877	18,843	16,461
UPZ O-D FE		X		X
Time of day Controls	X	X	X	X
Demographic Controls	X	X	X	X

Note: Data is from 1995 Mobility Survey. Low-Skill is a dummy for having no post-secondary education. Bus is a dummy for whether bus is used during a commute, relative to the omitted category of car. Data is from 1995. Time of day controls are dummies for hour of departure, and demographics are log age and a gender dummy. UPZ O-D FE are fixed effects for each upz origin-destination. Only trips to work included. Standard errors clustered at upz origin-destination pair. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 3: IV Results: Main Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	OLS	IV	IV	IV-LCP	IV-LCP	IV All	IV All
Panel A: Residents								
In(Res Floorspace Price)	0.421*** (0.077)	0.300*** (0.063)	0.165** (0.084)	0.274*** (0.068)	0.397*** (0.134)	0.476*** (0.104)	0.375*** (0.139)	0.471*** (0.108)
N	1,943	1,943	1,943	1,943	1,943	1,943	1,943	1,943
F-Stat					761.73	764.77	366.91	374.51
Over-ID p-value							0.69	0.12
In(Residential Pop)	0.238** (0.117)	0.248** (0.118)	0.168 (0.118)	0.174 (0.120)	0.279* (0.153)	0.299* (0.158)	0.260* (0.155)	0.281* (0.160)
N	1,997	1,997	1,997	1,997	1,997	1,997	1,997	1,997
F-Stat					1,687.78	1,643.78	828.30	811.19
Over-ID p-value							0.83	0.94
Panel B: Firms								
In(Comm Floorspace Price)	0.214** (0.108)	0.216* (0.113)	0.238** (0.110)	0.255** (0.115)	0.235 (0.147)	0.252* (0.154)	0.294** (0.147)	0.291* (0.154)
N	1,884	1,884	1,884	1,884	1,884	1,884	1,884	1,884
F-Stat					1,066.60	903.51	746.02	640.39
Over-ID p-value							0.27	0.48
Comm Floorspace Share	0.154*** (0.038)	0.152*** (0.040)	0.153*** (0.039)	0.154*** (0.041)	0.115** (0.054)	0.110* (0.059)	0.124** (0.055)	0.119** (0.059)
N	1,981	1,981	1,981	1,981	1,981	1,981	1,981	1,981
F-Stat					1,142.36	955.04	797.08	669.30
Over-ID p-value							0.44	0.51
In(Establishments)	1.165*** (0.325)	0.704** (0.314)	0.980*** (0.335)	0.636** (0.322)	1.359** (0.553)	1.300** (0.561)	1.000* (0.541)	0.829 (0.542)
N	1,724	1,724	1,724	1,724	1,724	1,724	1,724	1,724
F-Stat					186.13	215.47	291.39	246.75
Over-ID p-value							0.31	0.02
Locality Fixed Effects	X	X	X	X	X	X	X	X
CBD X Region Controls	X	X	X	X	X	X	X	X
Basic Tract Controls	X	X	X	X	X	X	X	X
Historical Controls		X		X		X		X
Init. Land Controls		X		X		X		X
Init. Demographic Controls		X		X		X		X
Distance to Tram Controls							X	X

Note: Observation is a census tract. Each entry reports the coefficient from a regression of the variable in each row on firm or residential commuter market access in first differences. Each column corresponds to a specification. Land market regressions use changes between 2000 and 2012, measuring the change in CMA induced by phases 1 and 2 of TransMilenio. Establishment regressions use changes between 2000 and 2015 and are weighted by number of establishments in 2000. Population regressions use changes between 1993 and 2005 measuring the change in CMA induced by phase 1 and are weighted by 1993 population. Only tracts further than 500m from a portal and the CBD (and less than 3km from a station) are included. 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 are log area and log distance to main road. Historical controls are dummies for quartile of 1918 population and a dummy for whether a tract is closer than 500m to main road in 1933. Initial land controls are the share of land developed, share of floorspace that is commercial, floor area ratio and log value of floorspace per square meter in 2000. Land market controls that represent initial values of outcome variable are excluded in each regression. Initial demographic controls are log population density and college share in 1993. Distance to tram is a dummy for whether a tract is closer than 500m from the historical tram line. Columns (1) and (2) run an OLS specification. Columns (3) and (4) instrument for the change in CMA holding residence and employment fixed at their initial levels and changing only commute costs, excluding the census tract itself from the variable construction. Kleinberg-Paap F-statistics are very high and not reported for brevity. Columns (5) and (6) instrument using the change in CMA induced by the LCP route, while (7) and (8) include both the LCP instrument and the change induced by the tram instrument. Standard errors clustered by tract reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 4: Commute Distance

Outcome	(1) OLS UseTM	(2) OLS lnDist	(3) IV lnDist	(4) IV All lnDist	(5) IV All lnDist
lnRCMA	0.957*** (0.212)	0.541** (0.252)	0.383 (0.300)	0.951** (0.404)	0.310 (0.466)
lnRCMA X High Skill					-0.147** (0.058)
N	9,088	22,119	22,119	17,212	19,920
R ²	0.07	0.10			
F-Stat				72.88	18.26
Over-ID p-value				0.52	0.55
UPZ FE	X	X	X	X	X
Locality FE X Post FE	X	X	X	X	X
Log Dist CBD X Region FE X Post FE	X	X	X	X	X
Trip Controls X Post FE	X	X	X	X	X
Tract Controls X Post FE	X	X	X	X	X
Historical Controls X Post FE	X	X	X	X	X
Educ X Post FE					X

Note: Observation is a trip, only trips to work are included. Column (1) reports coefficients from a regression of a dummy for whether an individual uses TransMilenio in 2015 on the change in lnRCMA in the origin UPZ. The other columns run difference-in-difference specifications using data from 2015 (Post) and 1995 (Pre), examining how changes in commute distances vary with changes in RCMA. RCMA is measured at the UPZ level using the pre-TM network in the pre-period and the 2006 network in the post period. Trip controls include hour of departure dummies and demographic characteristics (sex, log age, hh head dummy, occupation dummies). Tract controls include log area, log distance to a main road and log population density in 1993. Historical controls include quartile dummies of 1918 population, dummy for whether closer than 500m to main road in 1933, and when the tram instrument is used a dummy for whether a tract is closer than 500m from the historical tram line. Last column includes education level dummies interacted with Post FE. Standard errors clustered by origin UPZ are reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: College Share

Outcome: Change in College Share	(1) OLS	(2) OLS	(3) IV	(4) IV All
Δ lnRCMA	-0.011 (0.030)	-0.046 (0.029)	-0.040 (0.029)	-0.062 (0.047)
Δ lnRCMA X HighColl		0.051* (0.027)	0.063* (0.033)	0.111** (0.047)
N	1,886	1,886	1,886	1,886
R ²	0.27	0.27		
F-Stat				123.61
Over-ID p-value				0.55
Locality FE	X	X	X	X
HighColl FE		X	X	X
Log Dist CBD X Region FE	X	X	X	X
Tract Controls	X	X	X	X
Historical Controls	X	X	X	X

Note: Outcome is the change in a census tract's share of residents older than 20 with post-secondary education between 1993 and 2005. Dependent variable is change in RCMA between these years using the pre-TM and phase 1 of the system to measure commute times, interacted with a dummy for whether a tract is high college. The high college measure is constructed by first computing the share of college residents within a 1km disk around each tract centroid (excluding the tract itself) and then setting high college dummy equal to one for tracts in the top two terciles of its distribution. Specifications with interactions include an intercept to allow growth to differ across low and high college tracts (HighColl FE). Tract controls include log area, log distance to a main road and log population density in 1993; all other controls are as described in previous tables. Final column includes additional control for whether tract is closer than 500m from historical tram route. Columns (1) and (2) run OLS. Column (3) instruments for the change in CMA holding residence and employment fixed at their initial levels and changing only commute costs, excluding the census tract itself from the variable construction. Column (4) instruments using the change in CMA using the LCP and Tram instruments. Only tracts further than 500m from a portal and the CBD (and less than 3km from a station) are included. Robust standard errors reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 6: Wages

Outcome: lnWage	(1) OLS	(2) OLS	(3) IV	(4) IV-All	(5) IV-All
lnRCMA	0.479*** (0.162)	0.202* (0.108)	0.282** (0.129)	0.221 (0.221)	0.185 (0.236)
lnRCMA X College					0.298*** (0.054)
N	75,981	75,981	75,981	75,981	75,981
R ²	0.35	0.47	0.47	0.47	0.47
F-Stat				30.94	16.41
Over-ID p-value				0.94	0.64
UPZ FE	X	X	X	X	X
Region X Post FE	X	X	X	X	X
Log Dist CBD X Region FE X Post FE	X	X	X	X	X
Tract Controls X Post FE	X	X	X	X	X
Worker Controls X Post FE	X	X	X	X	X
College FE X Post FE		X	X	X	X
Historical Controls X Post FE				X	X

Note: Dependent variable is the log hourly wage for full-time workers reporting more than 40 hours worked per week. Data covers 2000-2005 in the pre-period and 2009-2014 in the post period. RCMA is measured at the UPZ-level using the pre-TM network in the pre-period, and using the 2006 network in the post-period. IV specification uses both the LCP and Tram instruments. Region are dummies for the North, West and South of the city. College is a dummy for having post-secondary education. Worker controls include gender and log age. Remaining controls are as described in previous tables. Standard errors are clustered by UPZ and period. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 7: Mode Choice Model Estimates

	(1)
Time	-0.012** (0.005)
Bus	-0.086* (0.050)
Car	0.837*** (0.292)
TM	-0.216** (0.105)
λ	0.140** (0.064)
Time of Day Controls	X
Demographic Controls	X

Note: Table shows estimation from nested logit regression on trip-level data from the 2015 Mobility Survey. λ is the correlation parameter for the public nest. Demographic controls include a sex dummy as well as dummies for quintiles of the age distribution, while time of day controls include dummies for the hour of trip departure. Each have choice-varying coefficients. Only trips during rush hour to and from work are included. Robust standard errors are reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 8: Gravity Regression

	(1)	(2)	(3)	(4)
High-Skill X In Commute Cost	-0.0242*** (0.0036)	-0.0240*** (0.0036)	-0.0239*** (0.0035)	-0.0234*** (0.0036)
Low-Skill X In Commute Cost	-0.0336*** (0.0043)	-0.0333*** (0.0043)	-0.0333*** (0.0042)	-0.0329*** (0.0042)
Crime		-0.005 (0.024)		-0.002 (0.019)
House Price		-0.318 (0.463)		-0.352 (0.353)
Primary Road		0.942 (0.624)		0.828 (0.580)
<i>N</i>	1,444	1,444	1,444	1,444
Origin-Skill-Car Ownership Fixed Effects	X	X	X	X
Destination-Skill Fixed Effects	X	X	X	X
Method	PPML	PPML	IV-PPML	IV-PPML

Note: Outcome is the conditional commuting shares between localities in 2015. Observation is an origin-destination-skill-car ownership cell. Skill corresponds to college or non-college educated workers. Only trips to work during rush hour (5-8am) by heads of households included. Columns 1 and 2 use PPML estimated under a GLM routine. Columns 3 and 4 implement IV-PPML with a 2-step GMM routine, using the times computed for both car and non-car owners under the LCP and Tram to instrument for times computed using the observed network. Crime, house price and primary road include the average number of crimes per year from 2007-2014, the average log house price in 2012, and the share of the trip that takes place along a primary road along the least-cost routes between origin and destination. In columns 1 and 2 standard errors are clustered by origin-destination locality; in columns 3 and 4 heteroscedasticity robust errors are recovered from the GMM variance matrix. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Table 9: GMM Results

Parameter	Estimate
Panel A: Firms	
μ_A	0.237*** (0.089)
Panel B: Workers	
η_L	3.595*** (0.861)
η_H	3.261*** (0.697)
μ_U^L	0.250*** (0.031)
μ_U^H	0.342*** (0.048)

Note: Estimates are from two-step GMM procedure separately for firms at the tract-industry level with 6137 observations and for workers at the tract-group-car ownership level with 7036 observations. Controls include log distance to CBD interacted with region fixed effects, commercial floorspace share in 2000, and log population density and college share in 1993 for employment moment conditions. Spillover parameter estimates obtained via delta method: original parameter clusters $\eta_L \mu_U^L$ and $\eta_H \mu_U^H$ are 0.898 (0.222) and 1.114 (0.176) respectively. Only tracts within 3km of the network and those more than 500m from portals and the CBD are included. Standard errors clustered at the tract reported in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 10: Welfare Approximations

	Welfare Low	Welfare High	Inequality
Value of Travel Time Savings	7.101	5.362	-1.739
First Order Approximation (CMA)	1.514	1.358	-0.156
First Order Approximation (Model)	5.063	3.587	-1.477
General Equilibrium	3.711	3.955	0.244

Note: Table shows the absolute value of the percentage change in welfare from removing TransMilenio under different approaches. The third column shows the log change in welfare inequality, defined as the ratio of high- to low-skill welfare. The first line is the VTTS approach. The second and third are first order approximations using the CMA regression elasticities and applying the envelope theorem to the social planner problem respectively. The fourth line is the full general equilibrium response from the model with equal home ownership.

Table 11: Effect of Removing Phases 1 and 2 of TransMilenio

	No Spillovers	Spillovers
Panel A: Closed City		
GDP	3.119	3.918
Rents	3.285	3.721
Welfare Low	3.444	3.814
Welfare High	3.651	4.169
Inequality	0.215	0.369
Panel B: Open City		
GDP	10.347	15.596
Rents	13.145	16.275
Population Low	8.562	10.744
Population High	9.543	12.303
Relative Population	1.072	1.747

Note: Table shows the (negative of the) value of the percentage change in each variable from removing phases 1 and 2 of the TransMilenio network from the 2012 equilibrium, with and without spillovers.

Table 12: Welfare Effects of TransMilenio: Decomposing the Channels

Model	Case	Low	High	Ineq.
Same η, θ , Perf Sub	Partial Eqbm	4.902	4.056	-0.889
Same η, θ , Perf Sub	Choices Adj	5.601	4.322	-1.354
Same η, θ , Perf Sub	Choices+Rents Adj	4.380	3.496	-0.925
Same η, θ , Perf Sub	GE	4.408	3.562	-0.885
Diff θ , same η , Perf Sub	GE	4.219	4.042	-0.185
Diff θ, η , Perf Sub	GE	4.163	4.010	-0.160
Diff θ, η , Imperf Sub	GE	3.814	4.169	0.369

Note: Table shows the percentage welfare and inequality change from TransMilenio under alternative adjustment scenarios. Each entry is computed by first simulating the effect of removing TransMilenio and then adding it back in, and reports the absolute value of the percentage welfare change from moving from the TM to no TM equilibrium for comparability with previous tables. Rows 1 to 4 consider a simplified model where worker groups share the same value for η, θ and are perfect substitutes in production. Row 1 assumes all choices and prices are fixed. Row 2 allows all discrete choices to adjust but holds prices fixed. Row 3 additionally allows house prices to adjust, while row 4 incorporates full GE adjustment. Rows 5 to 7 then compute the full GE effects where η, θ and σ are slowly turned to their estimated or calibrated values.

Table 13: Cost vs. Benefits of TransMilenio

	Closed City		Open City	
	No Spillovers	Spillovers	No Spillovers	Spillovers
Panel A: Costs & Benefits				
NPV Increase GDP (mm)	57,359	72,052	190,282	286,812
Capital Costs (mm)	1,137	1,137	1,137	1,137
NPV Operating Costs (mm)	5,963	5,963	5,963	5,963
NPV Total Costs (mm)	7,101	7,101	7,101	7,101
NPV Net Increase GDP (mm)	50,258	64,952	183,181	279,711
% Net Increase GDP	2.73	3.53	9.96	15.21
Panel B: Land Value Capture				
LVC Band Revenue (mm)	93	113	240	315
As share of capital costs	8.18	9.91	21.08	27.72
LVC CMA Revenue (mm)	170	203	464	571
As share of capital costs	14.96	17.86	40.82	50.21

Note: All numbers in millions of 2016 USD. NPV calculated over a 50 year time horizon with a 5% discount rate. Each column describes to a different model. Row (1) reports the increase in NPV GDP from phases 1 and 2 of the TransMilenio network from the baseline equilibrium in 2012 (calculated as the fall in GDP from its removal). Row (2) reports the capital costs of constructing the system, averaging 12.23mm per km over 93km of lines. Row (3) reports the NPV of operating costs, defined as farebox revenue in 2012. Row (4) reports the NPV of total costs, while row (5) reports the difference between row (1) and row (4). Row (6) reports this difference as a percent of the NPV of GDP in 2012. Row (7) reports the government revenue from the distance band-based land value capture scheme as described in the text, while row (8) reports this as a percentage of capital costs. Rows (9) and (10) report the same figures for the commuter market access-based LVC scheme.

Table 14: Model Extensions

	Closed City		Open City	
	Low Skill	High Skill	Output	Output
Baseline	3.814	4.169	3.918	15.596
No Fixed Costs	3.864	4.107	3.922	20.484
Preference Shocks	4.620	4.788	1.184	20.796
Joint Decision	4.106	4.178	3.937	20.725
Domestic Services	3.746	4.066	3.861	15.231
Home Ownership	3.758	3.851	4.224	17.337

Note: Table shows the (negative of the) value of the percentage change in welfare from removing phases 1 and 2 of the TransMilenio network from the 2012 equilibrium across different models. For each model, columns 1 and 2 report the percentage change in low- and high-skill worker welfare, column 3 reports the percentage change in output in the closed city model which column 4 reports the change in output in the open city model. Row 1 reports results from the baseline model. Row 2 presents those from the model without fixed costs (i.e. $\bar{h} = p_a = 0$). Row 3 reports the model with preference rather than productivity shocks by location of employment. Row 4 presents the model where there is a joint decision for residence and employment locations. Row 5 reports results from the model with employment as domestic servants. Row 6 shows results for the model with home ownership.

Table 15: Effect of Different System Components

	Welfare Low	Welfare High	Output
Add Cable Car	0.064	0.057	0.033
Remove Line South	1.609	1.560	1.045
Remove Line North	1.540	1.666	1.698
Remove Feeders	1.864	1.933	1.585
Uniform Reduction	2.993	3.146	3.716

Note: Table shows the (negative of the) value of the percentage change in welfare from removing a piece of the TransMilenio (existing, future or hypothetical) network. These counterfactuals are adding the Cable Car system, removing line H in the south, removing line A in the north, removing the feeder system, and a uniform 5.8% reduction in all commute times.

Table 16: Effect of Adjusting Housing Supply, and Land Value Capture Scheme

Panel A: Closed City						
	Output	Housing	Rents	Welfare Low	Welfare High	Gvt Revenue
Fixed Supply	3.918		3.721	3.814	4.169	
Free Adjustment	4.264	2.169	1.174	4.817	5.013	
LVC, Bands	4.171	1.007	2.731	4.073	4.462	9.91
LVC, CMA	4.218	1.977	1.396	4.765	4.954	17.86
Panel B: Open City						
	Output	Housing	Rents	Pop. Low	Pop. High	Gvt Revenue
Fixed Supply	15.596		16.275	10.744	12.303	
Free Adjustment	28.900	12.431	16.936	22.433	23.665	
LVC, Bands	18.279	2.792	16.233	12.654	14.545	27.72
LVC, CMA	21.585	5.128	17.038	16.118	17.412	50.21

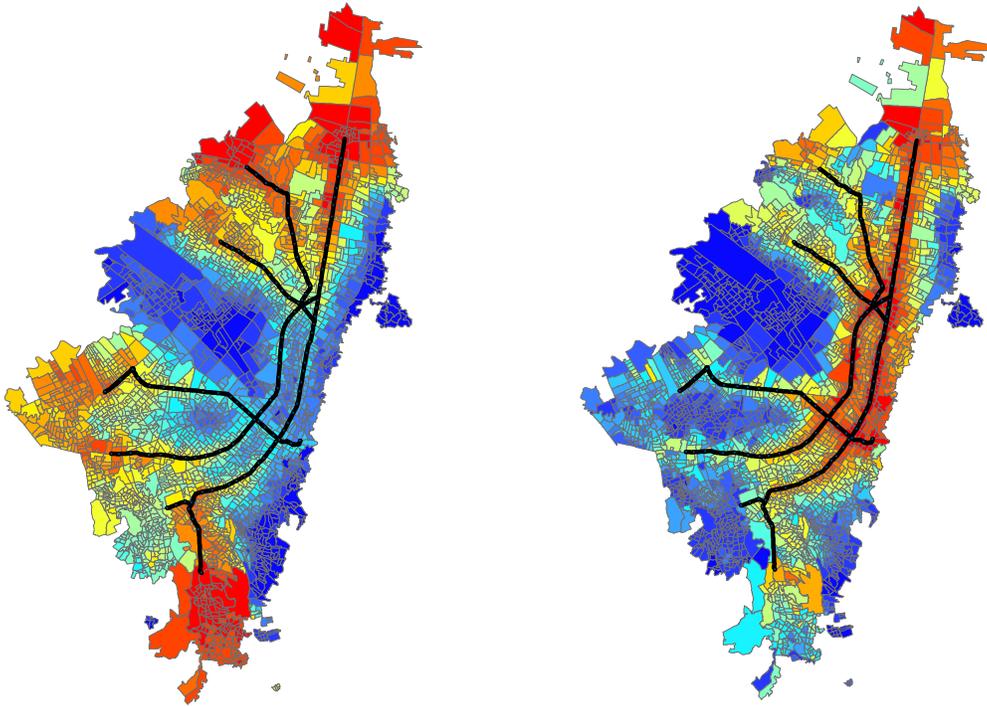
Note: Table shows the percentage change in each outcome going from the equilibrium without TransMilenio to that with TransMilenio under the housing supply conditions indicated in each row. Row (1) is the case with fixed housing supply. Row (2) is the case of freely adjusting housing supply. Row (3) is 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. Row (4) shows the results of the scheme based on predicted changes in commuter market access as described in the text. Government revenue from the scheme is given in column (6) as a percentage of construction costs.

Figures

Figure 1: Change in Commuter Market Access from TransMilenio

(a) Resident CMA

(b) Firm CMA



Note: Plot shows the baseline instrument for the change in CMA induced by holding population and employment fixed at their initial level and changing only commute costs. Tracts are grouped into deciles based on the the change in CMA, with warmer colors indicating a larger increase in CMA. Black line shows the TransMilenio routes as of 2006. See Section 7 for full discussion.

Figure 2: Population Density and Demographic Composition in 1993

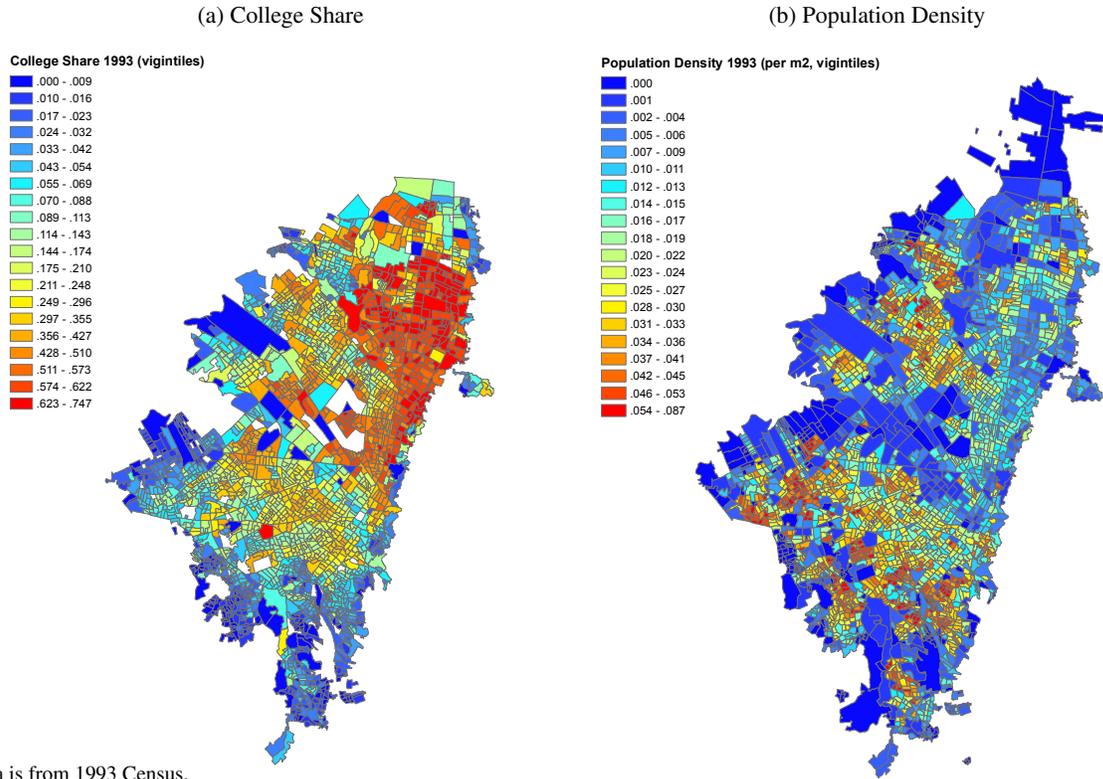


Figure 3: Employment Density and Industry Composition in 1990

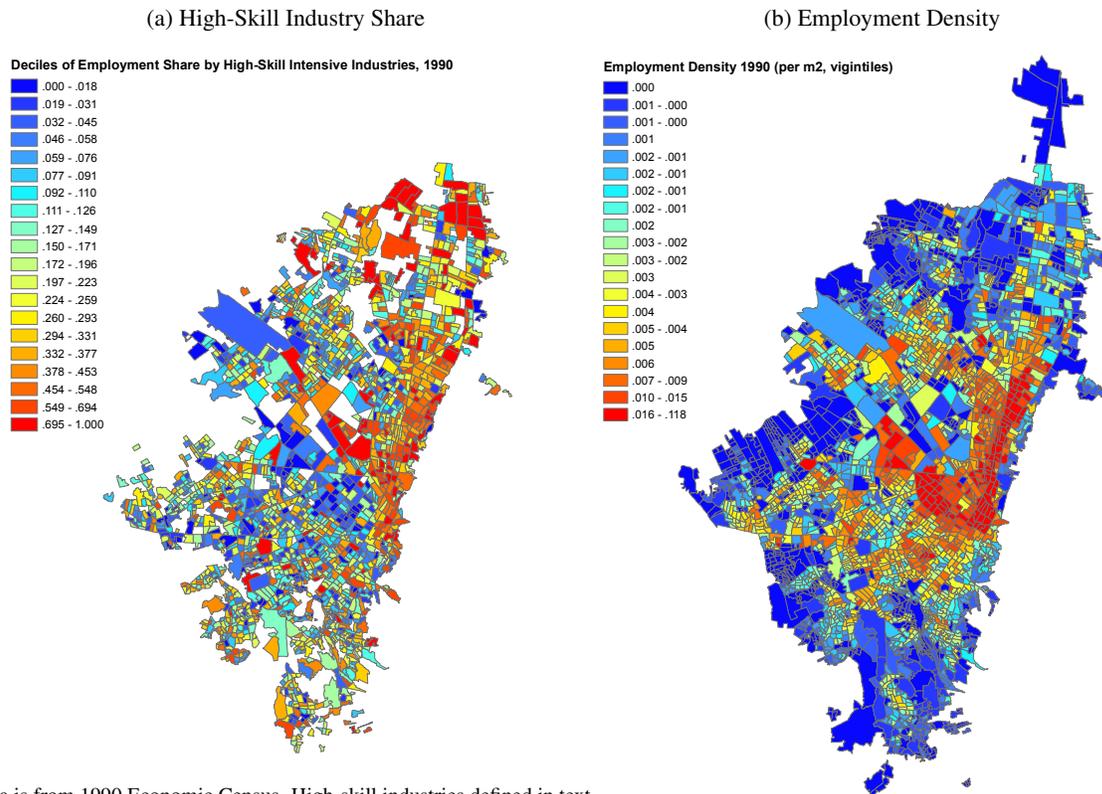
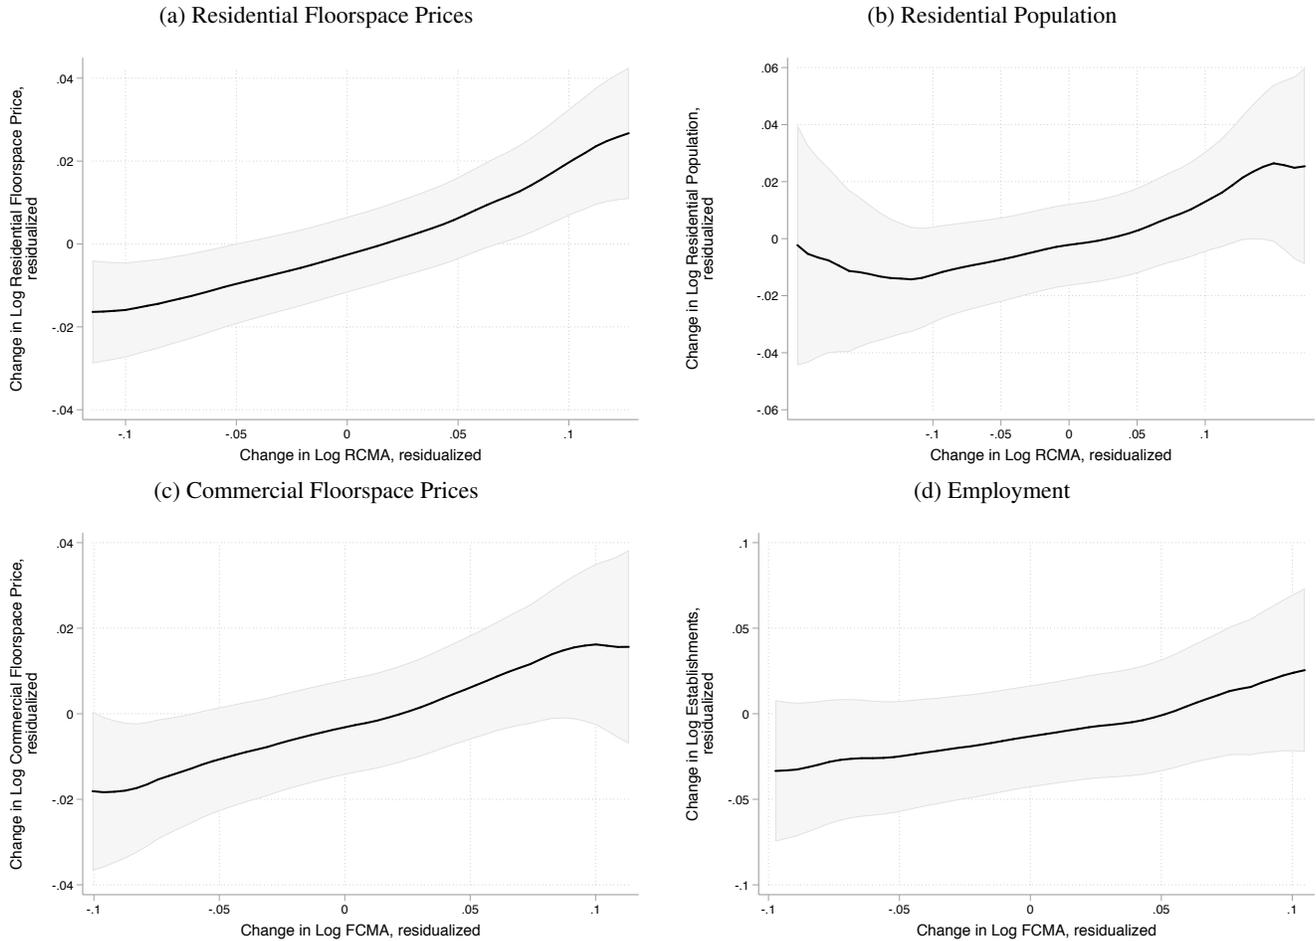
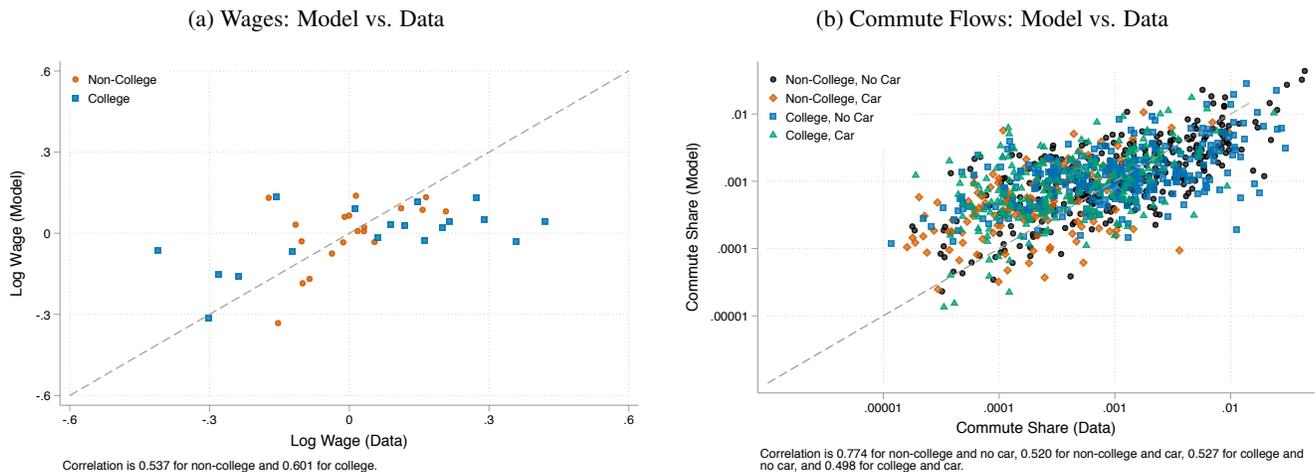


Figure 4: Non-Parametric Relationship Between Outcomes and Commuter Market Access



Note: Plot shows the non-parametric relationship between outcomes and CMA. Specifications correspond to the reduced form from column (4) of main table in which CMA is measured holding population and employment fixed at their initial levels, with the full set of baseline controls included, and is regressed directly on outcomes.

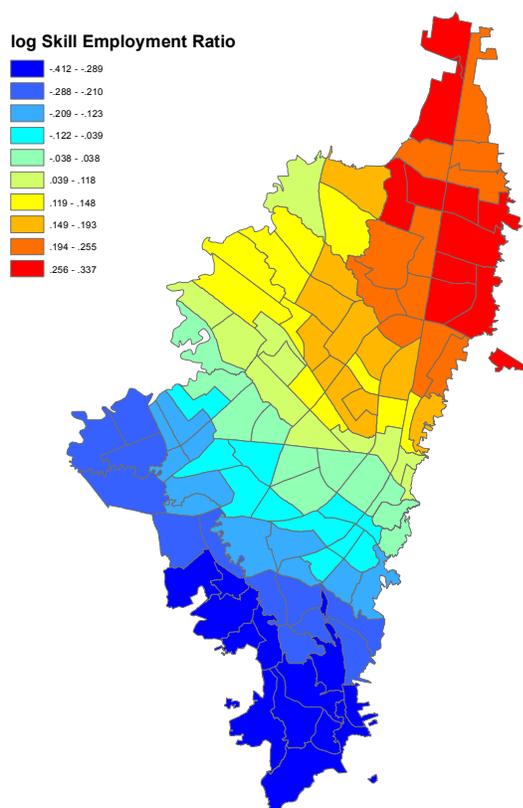
Figure 5: Wages and Commute Flows: Model vs. Data



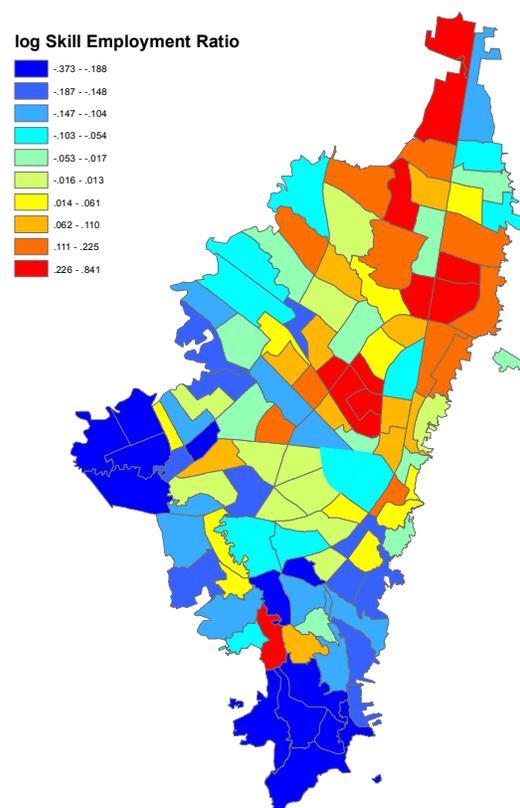
Note: Panel (a) compares the average wage by skill group in each locality as predicted by the model with that observed in the GEIH data (not used in estimation). In panel (b), a observation is a locality origin-destination pair, skill group and car ownership combination. Plot shows relationship between share of commuters choosing each (i, j, a) pair in the model vs those doing so in the 2015 Mobility Survey.

Figure 6: Relative Employment by Skill by UPZ: Model vs Data

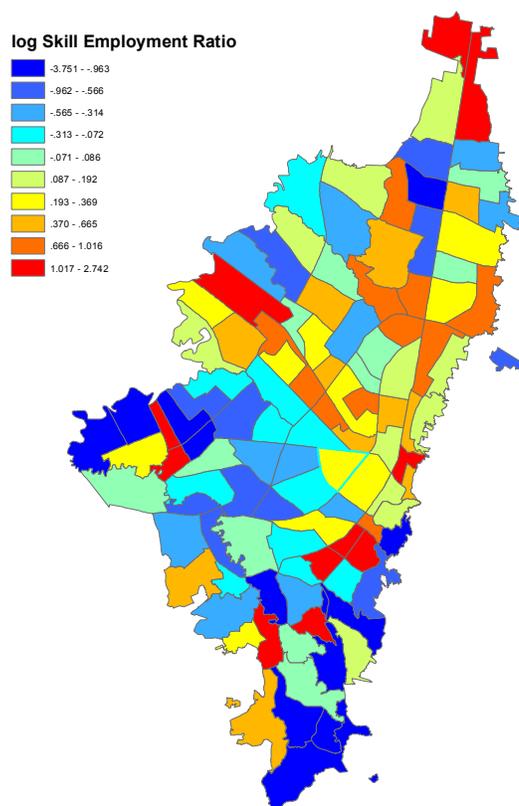
(a) Model: Perfect Substitutes & Same θ



(b) Model: Baseline Estimates

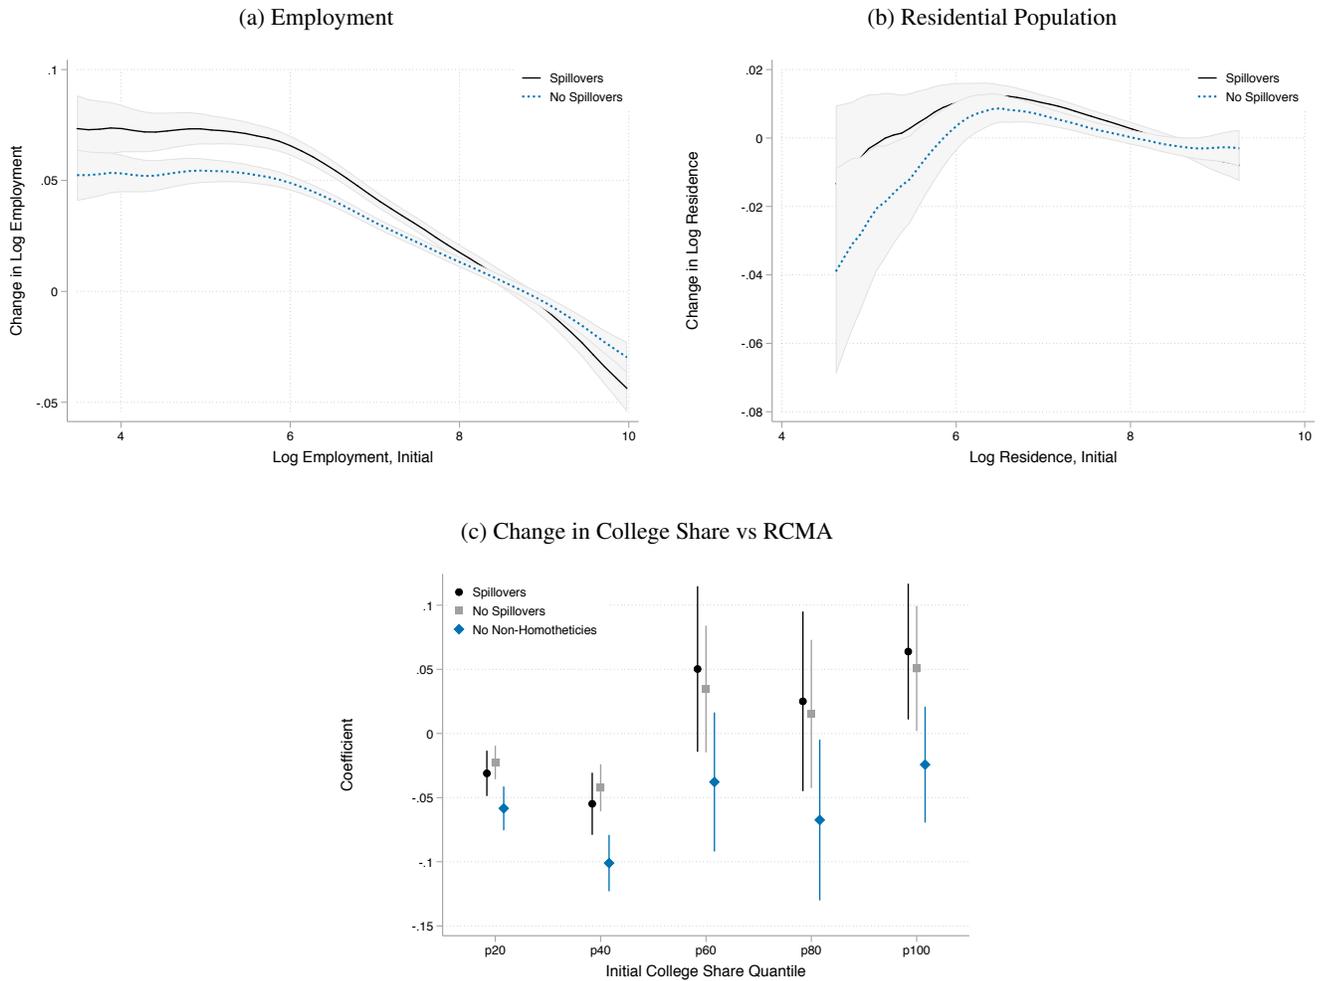


(c) Data



Note: Panel (a) shows the deciles of the distribution of the log skill employment ratio $\ln L_{FjH}/L_{FjL}$ by UPZ in the model when skill groups are perfect substitutes in production and have the same value of θ (equal to the average θ in the population). Panel (b) shows the distribution for the baseline model. Panel (c) shows the distribution in the 2015 Mobility Survey. Correlation between data in panel (a) and (c) is 0.256, while that between panel (b) and (c) is 0.406.

Figure 7: Simulated Changes in Outcomes



Note: Panels (a) and (b) plot the change in employment and population in each tract when TransMilenio is removed by each variable's initial level in the equilibrium with the system. Panel (c) plots the interaction terms from a regression of the change in college share on a full interaction between a dummy for a tract's initial college share quantile and the change in log RCMA in the baseline model with spillovers, without spillovers, and with $p_a = \bar{h} = 0$. The change in log RCMA is the same change in the average measure across groups as used in the empirical results, just now with counterfactual non-TM data generated by the model.