Subways and Urban Growth: Evidence from Earth*

Marco Gonzalez-Navarro
University of Toronto †
Matthew A. Turner
Brown University and *University of Toronto* ‡

12 December 2014

Abstract: We investigate the relationship between the extent of a city's subway network, its population and its spatial configuration. To accomplish this investigation, for a sample of the 632 largest cities in the world, we construct panel data describing the extent of each of the 138 subway systems in these cities, their population, and measures of centralization calculated from lights at night data. These data indicate that large cities are more likely to have subways, but that subways have at most a small effect on urban population growth. Consistent with economic theory and with other studies of the effects of transportation improvements on cities, our data also indicate that subways cause cities to be more decentralized.

Key words: subways, public transit, urban growth, urban decentralization.

JEL classification: L91, R4, R11, R14

*We are grateful to Fern Ramoutar, Mohamed Salat, Mahdy Saddradini and Farhan Yahya for their assistance compiling the subway data. We are also grateful to seminar participants at Brown University, LSE and PERC, and to Victor Aguirragabiria, Dwayne Benjamin, Gilles Duranton, Frank Kleibergen and Joan Monras for helpful comments and conversations. Turner gratefully acknowledges financial support from SSHRC and the financial support and hospitality of the Property and Environment Research Center, and the Enaudi Institute of Economics and Finance. Gonzalez-Navarro is grateful for financial support from the World Bank Urbanisation in Developing Countries Programme and the Ontario Work Study Plan. Both authors are grateful for financial support from the Societe du Grand Paris.

†121 St. George Street Toronto, ON M5S 2E8, Canada. email: marco.gonzalez.navarro@utoronto.ca. tel: +1(416) 9785-692.

‡Department of Economics, Box B, Brown University, Providence, RI 02912. email: matthew_turner@brown.edu.

1. Introduction

We investigate the relationship between the extent of a city's subway network, its population and its spatial configuration. To accomplish this investigation, for a sample of the 632 largest cities in the world, we construct panel data describing the extent of each of the 138 subway systems in these cities, their population, and measures of centralization calculated from lights at night data. These data indicate that large cities are more likely to have subways, but that subways have at most a small effect on urban population growth. Consistent with economic theory and with other studies of the effects of transportation improvements on cities, our data also indicate that subways cause cities to be more decentralized.

Our investigation is important for three reasons. First, understanding the effect of subways on cities is important if we are to evaluate proposals to build or extend subway systems. In 2010 Earth was home to 7,886 subway stations and about 10,700km of subway routes spread across 138 cities. Subway construction and expansion projects range from merely expensive to truly breathtaking. Among the 16 subway systems examined by Baum-Snow and Kahn (2005), construction costs range from about 25m to 550m USD2005 per km. On the basis of the mid-point of this range, 287m per km, construction costs for the current stock are about 3 trillion dollars. The Second Avenue subway extension, currently under construction in New York city, is estimated to cost almost 1.6b USD2010 per km, while Metropolitan Toronto is considering a 45b dollar(US) public transit expansion and Paris is considering a public transit expansion almost as large.

These costs are high enough that subway projects generally require large subsidies. To justify these subsidies, proponents often assert the ability of a subway system to have a transformative effect on the city and to encourage employment growth. A statement by the agency responsible for Toronto's transit expansion is typical: "Effective transit and transportation solutions can bolster our global competitiveness, protect our environment, and improve our quality of life. Expanding transportation can help create thousands of new green and well-paid jobs, and save billions of dollars in time, energy and other efficiencies."[1]

There is little evidence that subways have such transformative effects. Indeed, to date, there has been no city level statistical analysis of the effect of subway extent on any outcome other than ridership. If subways truly transform cities, as their proponents claim, then we should expect a migration of people into these cities. That our data do not provide evidence for such a migration suggests that the evaluation of prospective subway projects should rely more on the demand for mobility and less on the ability of subways to promote growth.

Understanding the effect of subways on cities is also important to policy makers interested in the process of urbanization in the developing world. Over the coming decades, we expect an enormous migration of rural population towards major urban areas. As a consequence, we can reasonably expect demands for urban infrastructure that are large relative to the ability of local

[1] http://www.metrolinx.com/en/regionalplanning/bigmove/big_move.aspx (accessed July 28, 2014).

and national governments to supply it. In order to assess trade-offs between different types of infrastructure in these cities, understanding the implications of each for welfare is clearly important. Since people move to more attractive places and away from less attractive ones (broadly defined), our investigation of the relationship between subways and population growth will help to inform these decisions. That subways have at most a tiny effect on population growth suggests that infrastructure spending plans in developing world cities should give serious consideration to non-subway infrastructure.

Finally, there is an active academic literature investigating the effect of transportation infrastructure on the growth and configuration of cities. In spite of their prominence in policy debates, subways have so far escaped the attention of this literature. This primarily reflects the relative rarity of subways. Most cities have roads so a single country can provide a large enough sample to analyze the effects of roads on cities. Subways are too rare for this. To conduct a statistical analysis of the effect of subways on cities requires data from, at least, several countries. An important contribution of this paper is to overcome this data problem by collecting data that describe all of the world's subway networks. In addition, with few exceptions, the current literature on the effects of infrastructure is static or considers panel data that is too short to investigate the dynamics of infrastructure's effects on cities. Because our panel spans the 60 year period from 1950 until 2010, we are able to investigate such dynamic responses to the provision of subways.

To estimate the causal effects of subways on urban growth and urban form, we must grapple with the fact that subway systems and stations are not constructed at random times and places. While much of the recent literature investigating the effects of infrastructure on cities relies on quasi-random variation to resolve this problem, a natural experiment to explain city level variation in subways seems unlikely. Instead, to solve this problem we exploit the panel nature of our data in two ways. First, our panel is long enough to allow us to estimate city specific intercepts and trends. If there are city specific levels or trends correlated with subways, then including city specific intercepts and trends can result in causal estimates.

However, it remains possible that subway expansions occur in non-random years, even after controlling for city effects. This might occur if, for example, cities build subways only after a sequence of years with improbably high growth or if there are short term population fluctuations as a consequence of subway construction activity. To resolve the inference problems created by this sort of equilibrium process for subway provision, we develop an econometric framework which models them explicitly. Practically, this leads us to instrument for current subways with long lags of subways. Implementing this econometric model does not substantively change the conclusions based on less sophisticated estimators.

2. Literature

A *Subways and urban growth*

With a few exceptions that we describe below, the literature that analyzes the effects of subways on cities consists entirely of analyses of a single city. Nevertheless, this literature is large and we here focus our attention on the small set of papers which attempt to resolve the problem of non-random assignment of subways. More complete surveys are available in Billings (2011) and Gibbons and Machin (2005).

Gibbons and Machin (2005) examine housing prices in London during the periods 1997-1999 and 2000-2001. These two year periods bracket two expansions of the London underground, six new stations and an extension of the Docklands light rail. Gibbons and Machin (2005) calculate various difference-in-differences estimates of the effect of these transit expansions on housing prices. They find that houses near a new transit station appreciate about 5% relative to houses further away, where a house within 2km is 'near' a subway station. Using a refinement of this estimator that allows distance to vary continuously, they find that moving one km away from a subway station decreases values by about 2% for the first two km, and about zero thereafter.

Billings (2011) conducts a similar exercise for a new light rail line in Charlotte, North Carolina. Charlotte opened a new light rail system in November 2007.[2] This system extended along one line, for about 15km, with 15 stations along the route. Like Gibbons and Machin (2005), Billings (2011) estimates the effect of subways on housing prices using a difference-in-differences estimator, where 'before' and 'after' refer to housing prices in the periods before and after the opening of the system. However, 'near' and 'far' are defined slightly differently. First, distance is defined as distance to the subway line rather than distance to a station. Given that stations are about 1km apart this is probably not important. Second, houses 'far' from the station are restricted to be close to alternate corridors that were candidates for a transit network that was ultimately not built, on the grounds that houses in alternate corridors are likely to resemble those in the successful corridor in unobserved ways. Despite the differences in milieu and method, Billings (2011) arrives at estimates quite close to those of Gibbons and Machin (2005): single family houses within 1.6km of the transit line see their prices increase by about 4% while condominiums see their prices rise by about 11%. Note that, like Gibbons and Machin (2005), Billings (2011) observes that changes result from subway construction over the course of just a few years.

Each of these papers makes a credible attempt to overcome the fact that subway systems are not located randomly within cities. However, neither provides us with much information about the relationship between subways and growth. If subways affect the growth of cities, then they may affect it everywhere, both near and far from the station. Such citywide effects are, by construction, invisible to the differences-in-differences methodology.

[2]The Charlotte light rail system is not completely isolated from pedestrian and automobile traffic and so does not appear in our data as a subway.

Therefore, while the existing literature makes some progress on the problem of non-random assignment of subways to places, it does so at a high cost. The difference-in-differences methodology cannot tell us about the effect of changes in the overall level of activity within a city and, unless we are specifically interested in reorganizing economic activity across neighborhoods within a city, it is such changes in the overall level which are of primary policy interest and which are the object of our investigation.

B *Other infrastructure and urban growth*

The literature relating highway and railroad infrastructure to urban growth has developed rapidly over the past several years and is surveyed in Redding and Turner (2015). For our purposes, it can be divided into studies of the effects of infrastructure on the internal structure of cities and studies of the evolution of city level population or employment.

Baum-Snow (2007) investigates the effect of radial highways on population decentralization for a sample of large US cities between 1950 and 1990. He relies on an early plan of the highway network that was drawn for military purposes as a source of quasi-random variation in highway placement. He finds that, over the whole 40 year course of his study period, a single radial highway causes about a 9% decrease in central city population. Importantly, he is able to measure partial effects over 20 year time periods. Baum-Snow, Brandt, Henderson, Turner, and Zhang (2012) conduct a similar exercise. For a sample of Chinese cities between 1990 and 2010, Baum-Snow *et al.* (2012) investigate the effect of highways and railroads on the decentralization of population, employment and production. They rely on historical networks predating economic reforms of the 1990s as a source of quasi-random variation in highway and rail placement. They find that radial highways contribute to the decentralization of population, that railroads contribute to the decentralization of production, and that ring road capacity contributes to the decentralization of both. Again, Baum-Snow *et al.* (2012) find measurable effects over the whole of their 20 year study period and also over ten year sub-intervals. Garcia-López (2012) examines the effect of radial highways on Spanish cities between 1992 and 2007 and also finds that highways decentralize population, again over a 10-20 year horizon.

Probably more relevant to the present investigation, for a sample of US cities between 1980 and 2000, Duranton and Turner (2012) use a research design similar to Baum-Snow (2007) to investigate the relationship between urban employment growth and a city's stock of interstate highways. They find that a 10% increase in highways in a city causes about a 1.5% increase in the population of a city over 20 years. Using a similar research design, Garcia-López, Holl, and Viladecans-Marsal (2013) finds that highways cause about the same rate of population growth in Spanish cities.

In addition to investigating the effects of highways using city level variation, both Baum-Snow (2007) and Garcia-López (2012) examine population density within cities as a function of distance to a highway. They find that population tends to be denser near highways and decreases with distance to a highway. Like highways, subways are intended to provide mobility for people and

the available evidence suggests that both highways and subways usually attract economic activity, broadly defined. Therefore, to the extent that subways affect urban form and growth, we can reasonably expect to observe these effects over a 10-20 year period. This is close to the duration of our panel for lights at night and is shorter than the period for which we observe population.

The relationship between highways and railroads and urban growth is better understood than is the relationship between subways and urban growth for two reasons. First, highways and railroads are pervasive and so even medium sized countries can provide statistically useful samples of cities. Second, the literature has devised credible instrumental variables strategies to deal with the non-random assignment of highways and railroads to cities. An analysis of subways, on the other hand, requires data from many countries to construct a reasonably large sample and there is little hope for quasi-random variation in the assignment of subways. Two important contributions of the present investigation are to overcome the data problem and to develop an identification strategy that relies on time series variation in panel data.

Subways and city level outcomes

The only studies (of which we are aware) to investigate the effects of subways on city level outcomes are primarily or completely interested in ridership.[3] On the basis of a single cross-section of about 50 cities, Gordon and Willson (1984) conduct a city level regression to predict riders per mile of track as a function of city population density and country level per capita GDP. They find that these two variables are excellent predictors of ridership. Barnes (2005) provides evidence from a few cities in the US for an unsurprising refinement of this intuition. People are much more likely to take transit for trips to a central business district than for trips to other locations. Finally, Baum-Snow and Kahn (2005) provide evidence from 16 US cities for a similar relationship between density and transit use, although their small sample size limits the precision of their results. They also show that ridership in catchment areas for new stations attains almost the same level as in the catchment areas of old stations over their 30 year study period. On a more cautionary note, and consistent with the finding in Gordon and Willson (1984) that ridership decreases with income and increases with density, Baum-Snow and Kahn (2005) find that most US transit expansions have only small effects on ridership, a conclusion echoed in Gomez-Ibanez (1996) for time series data on the use of Boston's transit system.

Aside from suggesting that pessimism about the growth of transit ridership is in order, these papers illustrate two points important to our analysis. First, they demonstrate the incompleteness of our understanding of the relationship between subways and the fortunes of cities. For practical purposes, there is no systematic evidence to support the claims of subway proponents that sub-

[3]We note the literature on modal choice using individual level data. This important literature is only tangentially related to our present inquiry. A survey is available in Small and Verhoef (2007).

Table 1: Descriptive statistics for the world's subway systems and cities in 2010

	World	Africa	Asia	Aus.	Europe	N. America	S. America
All cities							
N	632	73	341	6	57	99	56
Total Stations	7,886	51	2,977	0	2,782	1,598	478
Total route km	10,686	56	4,224	0	3,558	2,219	628
Population	2,427	2,104	2,511	2,429	1,921	2,441	2,825
log(Pop.)	14.329	14.296	14.325	14.572	14.227	14.344	14.450
Δ log(Pop.)	0.180	0.248	0.198	0.107	0.046	0.143	0.189
Δ_t^2 log(Pop.)	-0.010	-0.014	-0.008	-0.006	-0.005	-0.013	-0.015
Mean light in 25km disk	62	24	52	65	84	116	58
Corr. lights & pop.	0.57	0.56	0.58	0.85	0.69	0.71	0.92
Subway cities in 2010							
N	138	1	53	0	40	30	14
Stations	57	51	56		70	53	34
Δ Stations	3.758	4.250	4.568		3.965	2.658	2.423
Δ_t^2 Stations	0.691	0.000	1.374		0.325	0.115	0.435
log(Stations)	3.551	3.932	3.505		3.872	3.332	3.252
Δ_t^2 log(Stations)	0.204	0.101	0.305		0.158	0.140	0.117
Δ_t^2 log(Stations)	-0.047	0.000	-0.054		-0.018	-0.066	-0.088
Route km	77	56	80		89	74	45
Pop. per route km	61	196	75		25	65	141
Pop. per station	82	216	106		32	90	185
Pop.	4,706	11,031	5,951		2260	4,814	6,300
log(Pop.)	14.934	16.216	15.153		14.380	15.051	15.344
Δ log(Pop.)	0.113	0.124	0.144		0.045	0.123	0.170
Δ_t^2 log(Pop.)	-0.011	-0.014	-0.012		-0.005	-0.013	-0.017
Mean light in 25km disk	122	212	117		95	171	109
Corr. lights & pop.	0.67		0.67		0.69	0.78	0.91

Note: Population levels reported in thousands. Lights data are based on radiance calibrated lights at night imagery. All entries describing levels report 2010 values. Entries describing changes are averages over the period from 1950 to 2010.

ways transform cities. Second, they point to the importance of putting subways in dense places and of allowing dense development in areas served by subways.

3. Data

To investigate the effect of subways on the evolution of a city's population and its spatial structure we require data describing subways, population and spatial structure for a panel of cities. We construct such data from three principal sources. Our population data are the UN World Cities Data, our subway data are the result of primary data collection and our description of urban spatial structure derives from lights at night data.

In addition to our three main data sources, for each of the cities in our sample we construct several time-invariant control variables. They are continent of city, capital city status, distance to the ocean, distance to nearest land border and distance to the nearest major navigable river. The UN World Cities Data reports continent and capital city status. Using GIS software, we calculate country, distance to the ocean, to the nearest navigable river and to the nearest land border on the basis of shapefiles describing oceans, rivers and international boundaries. We also use the Penn World Table measure of country level GDP per capita as a time varying control in many of our regressions.

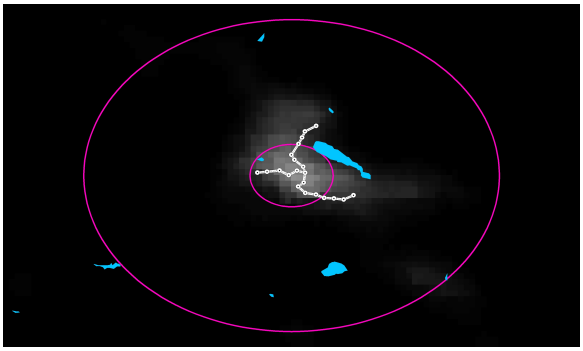
A *Population data*

Our data are organized around the UN World Cities Data.[4] Produced by the United Nations, Department of Economic and Social Affairs, Population Division, these data describe population counts for all cities whose population exceeds 750,000 at any time during 1950-2010.

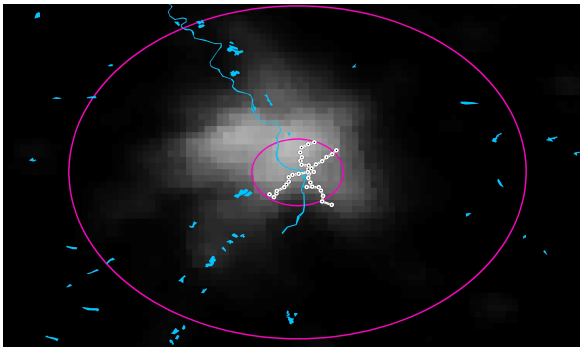
Constructing international data describing city level population is subject to two problems, the availability of population data and cross-country differences in the definitions of cities. Population data are generally available from decennial or quinquennial censuses. However, census years do not synchronize neatly across countries. To resolve this problem, the UN World Cities Data interpolate to construct annual values. Therefore, because few countries conduct censuses more often than every five years, successive annual population changes must sometimes reflect linear interpolation of the same proximate census years. To avoid making inferences from such imputed population changes, we restrict attention to observations drawn every fifth year, e.g., 1950, 1955, ..., and refer to each such observation as a 'city-year'. This decreases the likelihood that sequential city-years are calculated by interpolation from the same two underlying censuses. Resolving the problem of differential city definitions relies on the judgement of the UN researchers who assembled the data. In private correspondence, these researchers assured us that they understood

[4]Downloaded from `http://esa.un.org/unup/GIS-Files/gis_1.htm`, February 2013.

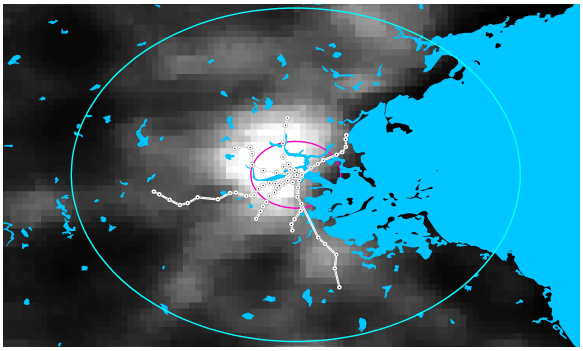
Figure 1: Lights and subways in 2010 for six cities



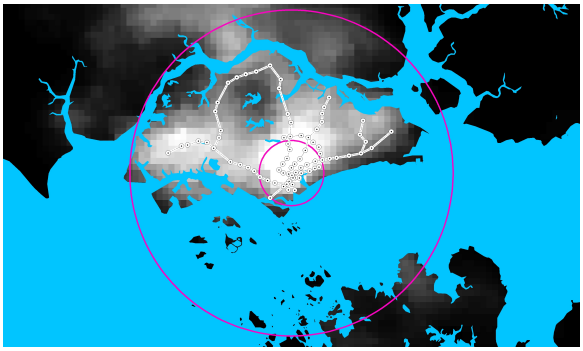
Tibilisi: 1.1m pop, 21 stations



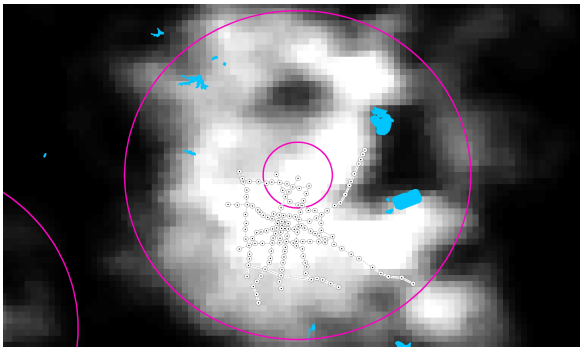
Toulouse: 0.9m pop and 37 stations



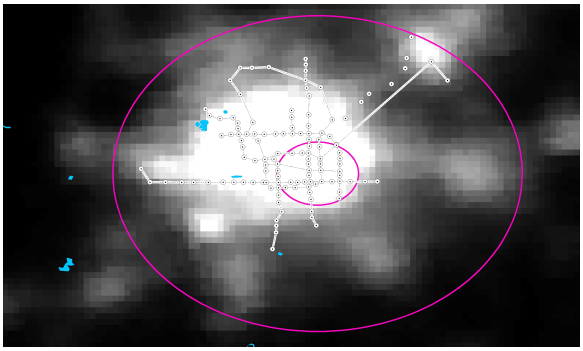
Boston: 4.7m pop, 74 stations



Singapore: 5.1m pop, 78 stations



Mexico City: 20.1m pop and 147 stations



Beijing: 15m pop and 124 stations

Note: Images show 2010 radiance calibrated lights at night, 2010 subway route maps and all subway stations constructed prior to 2013. The gray/pink ellipses in each figure are projected 5km and 25km radius circles to show scale and light gray/blue is water.

this problem and endeavored to make their 'cities' as consistent as possible across countries and time.

The top panel of Table 1 describes our population data. Our sample consists of 632 cities, more than half in Asia. In 2010, the mean size of a city in our sample is about 2.4 million. There is little variation in mean size across continents, although cities in South America tend to be larger while cities in Europe tend to be smaller. Between 1950 and 2010, the mean five year growth rate of a city in our sample is about 18%. This rate falls by about 1% every five years. Not surprisingly, cities in Africa, Asia and South America grow faster than in North America and Europe. Europeans cities are the obvious outlier and grow more slowly than cities elsewhere. The growth rate of cities is declining on all continents and this decline is somewhat slower in Europe.

The bottom panel of Table 1 describes our population data for the 138 cities in our sample with a subway in 2010. At 4.7m people on average, these cities are dramatically larger than non-subway cities. Cairo is the single African city with a subway, and so the Africa column in the bottom panel of table 1 is really a 'Cairo column'. Australia has no subway cities in 2010. Asian and South American subway cities are larger than those in North America and dramatically larger than those in Europe. The growth rate for an average subway city is about 13%, somewhat slower than in the whole sample. As for the whole sample, European subway cities are growing more slowly than other subway cities. Also similar to the whole population of cities, growth rates are declining by about 1% every five years and this decline is somewhat slower in Europe.

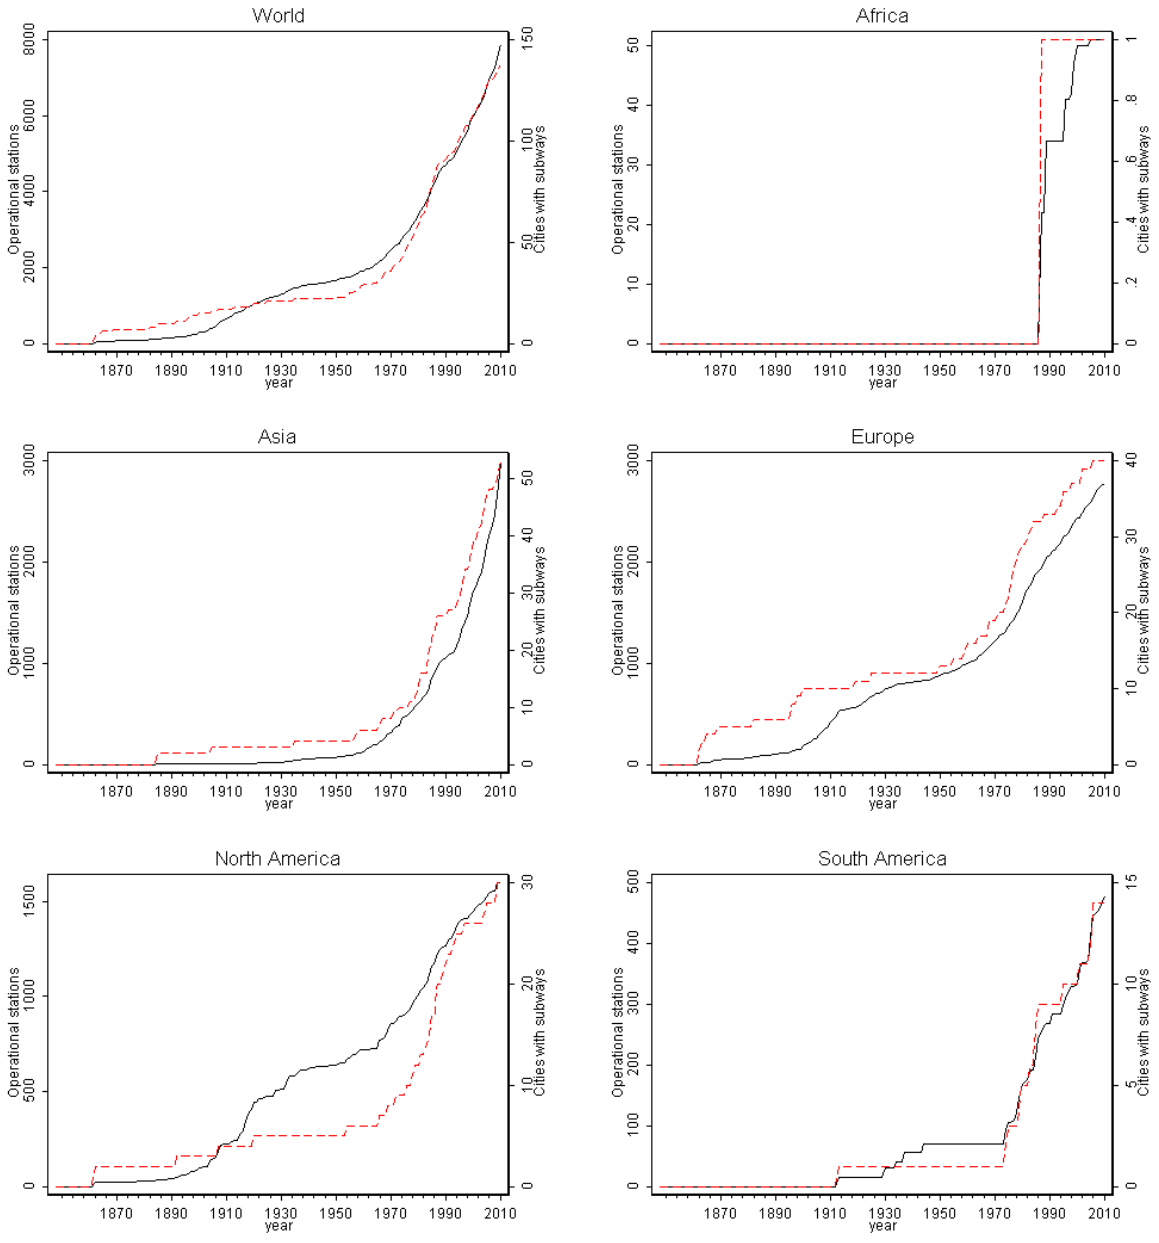
B *Subways data*

Our data describe the latitude, longitude and date of opening of every subway station in the world. These data were compiled manually between January 2012 and February 2014 using the following process. First, using online sources such as `http://www.urbanrail.net/` and links therein, together with links on wikipedia, we complied a list of all subway stations worldwide. Next, for each station on our list, we record opening date, station name, line name, terminal station indicator, transfer station indicator, city and country. This process leads us to enumerate subway stations in 161 cities. Of these, 138 are large enough to appear in the UN World Cities Data and are the main subject of our analysis.[5]

For our purposes, a 'subway' is defined as an electric powered urban rail that is completely isolated from interactions with automobile traffic and pedestrians. This excludes most streetcars, because they interact with vehicle traffic at stoplights and crossings, although we include underground streetcar segments. In order to focus on intra-urban subway transportation systems, we also exclude heavy rail commuter lines. We do not distinguish between surface, underground or

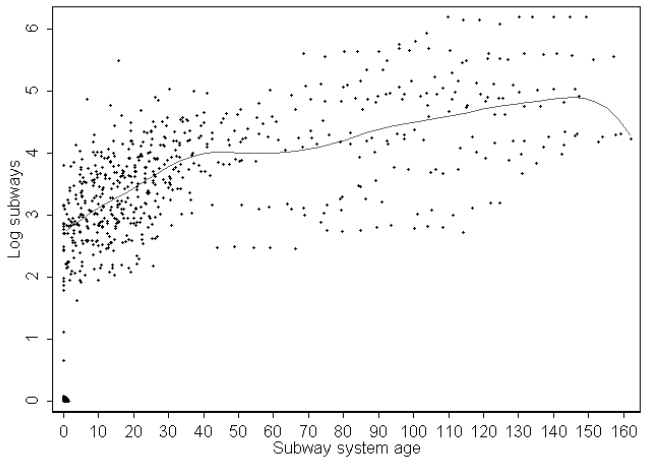
[5]The 23 cities with subways in 2010 that do not occur in our population data are: Bielefeld, Bilbao, Bochum, Catania, Dortmund, Duisburg, Dusseldorf, Essen, Frankfurt, Genova, Hannover, Kitakyushu, Kryvyrih, Lausanne, Mulheim, Naha, Nuremberg, Palma, Perugia, Rennes, Rouen, Seville and Wuppertal.

Figure 2: Growth of subway systems for the whole world and by continent



Note: In each graph, the dashed line indicates the number of cities with a subway system and the solid line indicates the total number of operational stations. Australia has no subway systems and is omitted.

Figure 3: Stations in a subway system by time since system opening



Note: Vertical axis is log of operation stations in a system. Horizontal axis is years since station opening. Dots indicate individual city-years.

aboveground subway lines as long as exclusive right of way condition is satisfied. These subways systems typically operate frequently, e.g., 10 minute headways or less during daytimes, are quick, reliable and are used mostly for the intra-urban transportation of people. For the most part, our subways data describe public transit systems that would ordinarily be described as 'subways', e.g., the Paris metro and the New York city subway, and only such systems. As with any such definition, the inclusion or exclusions of particular marginal cases in our sample may be controversial.

We use our data to construct three measures of subway extent for each city-year. Most simply, we count the number of operational stations in each year. Since our data also enumerate subway lines, we also count the number of operational subway lines in each city in each year. Finally, by connecting stations on each subway line by the shortest possible route, we approximate the route of each subway line. Taking the union of all such lines in a city approximates each city's network. Calculating the length of this network gives us the length of each system. In this way we arrive at our three primary measures of subway extent for each city-year; operational stations, operational lines and route kilometers.

Figure 1 illustrates our subway data for six cities. The figure shows all stations operational prior to 2013 as black dots. The network maps, on which the 2010 calculation of route km is based, are shown as black lines. Stations that opened after 2010 are not connected to the network. In each panel of the figure, a gray circle or ellipse describes a circle of 25km radius to show scale. That this circle is distorted in Northerly cities is a consequence of our map projection. To show the configuration of each city, the background shows lights at night in 2010. In the top row, with 2010 populations of 1.1m and 0.9m Tibilsi (Georgia) and Toulouse (France) are among the smallest cities in our sample to have subways. In 2010 their subway systems consist of 21 and 37 stations, and

27 and 28 route km. In the middle row, Boston and Singapore have populations of 4.7m and 5.1m, near the 4.7m mean for subway cities. Their subway systems consist of 74 and 78 stations and of 88 and 111 route km, which makes both systems somewhat larger than both world and the relevant continental averages. The bottom row of figure 1 shows two of the largest cities in our sample, Mexico City and Beijing. The population of Mexico City in 2010 was just over 20m against about 15m for Beijing. Their subway systems contained 147 and 124 stations and consisted of 182 and 209 route kilometers.

Figure 1 reveals two noteworthy features of our subways data. First, in each of the six cities, only a small portion of the city is within walking distance of a subway and the catchment area of the subway is centrally located. This is typical. An average city in our sample has about 57 stations. Of these, about 9% are within 1500m of the center, about 29% are between 1500m and 5km of the center, about the same share lie between 5 and 10km and between 10 and 25km. Just 7% of stations are beyond 25km from the center. Since the area to be served expands quadratically, this means that subways per square kilometer decreases rapidly in successive annuluses. In an average subway city, there are 0.67 stations per km^2 within 1500m of the center, 0.22 stations per km^2 between 1500m and 5km from the center, 0.07 stations per km^2 between 5 and 10km from the center, and 0.001 stations per km^2 between 10 and 25km from the center. Thus, in an average city, the preponderance of the subway system is located within 10km of the center, and much of it is even closer. Station density declines rapidly with distance from the center.

Close inspection of the network maps in figure 1 suggests that our networks probably diverge slightly from the actual network. The algorithm that we use to construct network maps connects all open stations on a subway line by the shortest possible route. This can result in divergence from the actual route for four reasons. First, if pairs of stations are connected with curving track, the actual route will diverge from our straightline network. Second, if intermediate stations on a line open after the end points, then the algorithm will not include the intermediate stations on the network until they open. Third, we may mis-attribute stations to subway lines. Fourth, if a route is served by two or more sets of tracks, then this replication is invisible to us. Loosely, our measure of length is a measure of the route kilometers required to serve operational stations in each year rather than a literal measure of the length of track in the system. While we regard the route kilometers measure as being of considerable interest, we suspect it is a noisier measure of subway extent than is the count of operational stations. Given this, our investigation relies primarily on the count of operational stations to measure system extent.

Table 1 presents descriptive statistics for the world's subway systems in 2010. In 2010 there were 7,886 operational subway stations and 10,686 route kilometers of subways in the world, divided across 138 operational systems. Of these 138 subway cities, 53 are in Asia, 40 in Europe, 30 in North America, 14 in South America, one in Africa and none in Australia. Asia, Europe, North America and South America account for 38, 35, 20 and 6 percent of all operational stations in 2010. The corresponding percentages of route kilometers are 40 for Asia, 33 for Europe, 21 for North

America and 6 for South America. Thus, Asia has more systems than Europe, but a typical system in Europe has more stations and route kilometers. North America accounts for a small share of subway stations and route km, it contains a small number of systems and the average extent of these systems is, on average, between that of Asian and European systems.

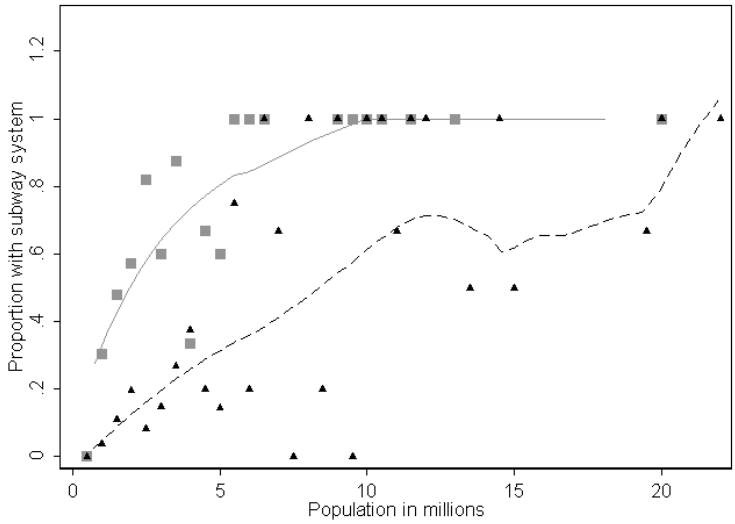
Table 1 reveals huge differences in the availability of subways across continents. Of the 341 large cities in Asia only 53, about 15%, have subway systems. In Europe, more than two thirds of large cities have subways, while in North America it is just less than one third. South America is a bit lower at 25%. Conditional on being in a subway city, the level of service also varies widely by continent. Cities are smaller and subway systems larger in Europe where there are 25,000 people per route km and 32,000 per station. These service levels are higher than those in North America and Asia and higher still than those in African and South American subway cities. Interestingly, although the share of North American cities with subways is much higher than in Asia, people per station and people per route km in subway cities are close for the two continents.

Two features of table 1 stand out. First, the huge gap in subway provision between Europe and the rest of the world. Second, the weak connection between mean city size and subway extent. In particular, Asia is home to the preponderance of the world's large cities while South America's cities are larger, on average, than those elsewhere. However, neither South America nor Asia is well provided with subways relative to Europe and North America. Indeed, Europe's cities are the smallest and it is by far the best provided with subways.

Table 1 provides a snapshot of the world's subway systems and cities in 2010. Figure 2 illustrates the expansion of the world's subway systems over the past century, both in aggregate and by continent. There were five subway systems in operation prior to or during 1860; Liverpool, Boston, London and New York. The "L" opened in Chicago in 1892 and The Paris Metro opened in 1900. Both the aggregate world data and the continental data, except for Asia, show a first wave of subway construction between the two world wars and a second wave beginning in the 1970s and continuing to 2010. The growth of Asian subways begins in the 1970s and has accelerated since. Except for North America, expansion of subway systems and increases in the number of subway cities track each other closely. The divergence between subway cities and stations in North America reflects the growth of the New York city subway system. In 2010, the 1169 subway stations operating in the US were spread across 21 cities. However, 489 of these stations were in New York. Chicago is the second largest system at 142 stations. On average, the remaining 19 US subway cities have just 29 stations each.

Where figure 2 shows the growth of the world's subways, figure 3 traces out the size of individual systems as a function of the time since they opened. Each marker in this figure describes a city year, so that there is one marker for each of the city-years in our data where at least one subway station is open. Consistent with figure 2, most of the observations are in the left portion of the graph. This reflects the fact that many subways systems have opened in the past 30 years. On the other hand, markers in the right hand portion of the graph describe the handful of subway

Figure 4: Proportion of cities with subways systems by population for two income classes



Note: Gray squares correspond to rich country cities and black triangles to poor country cities. See footnote 10 for the list of countries.

systems that date back to the 19th century. The solid line in the figure describes a locally weighted regression of system size on system age. This figure suggests that the expansion of a city's subway network is predictable. Expansion is rapid and approximately loglinear during the first 30-40 years after a system opens. After a system is about 40 years old, growth slows but remains approximately log linear, though at a lower growth rate.

C *Lights data*

Lights at night data are collected by earth observing satellites that measure the intensity of visible light every night in 30 arc second cells, about one kilometer square, on a regular grid covering the entire world. Most extant applications of the lights at night data in economics rely on the "DMSP-OLS Nighttime Lights Time Series".[6] These data are available annually from 1992 until 2012. Each of these lights at night images is a composite constructed from many raw satellite images and the value for each cell reflects average light intensity, over all cloud free images, on a scale of 0-62 with 63 used as a topcode. Since most large cities, particularly in the developed world contain large topcoded regions near their centers, these data are of limited use for studying the internal structure of the large wealthy cities where most subways are located. We exploit 'radiance calibrated at night data',[7] collected during times when the satellite sensor was set to be less sensitive. These data are less able to distinguish dim light sources, but are able to measure variation in light within regions

[6] Available from `http://www.ngdc.noaa.gov/dmsp/downloadV4composites.html` (October 2014). We are grateful to Alexi Abrahms for drawing our attention to these data.

[7] Downloaded in October 2014 from `http://ngdc.noaa.gov/eog/dmsp/download_radcal.html`.

Table 2: Population and subway stations for the world's 90 largest cities as of 2010.

City Name	Pop.	Stations	Stations pp.	City Name	Pop.	Stations	Stations pp.
Tokyo	36,933	255	0.69	Ho Chi Minh City	6,189	.	.
Delhi	21,935	128	0.58	Miami	5,971	22	0.37
Mexico City	20,142	147	0.73	Santiago	5,959	93	1.56
New York	20,104	489	2.43	Baghdad	5,891	.	.
Sao Paulo	19,649	62	0.32	Philadelphia	5,841	64	1.10
Shanghai	19,554	239	1.22	Nanjing	5,665	54	0.95
Mumbai	19,422	.	.	Haerbin	5,496	.	.
Beijing	15,000	124	0.83	Barcelona	5,488	137	2.50
Dhaka	14,930	.	.	Toronto	5,485	69	1.26
Kolkata	14,283	23	0.16	Shenyang	5,469	22	0.40
Karachi	13,500	.	.	Belo Horizonte	5,407	19	0.35
Buenos Aires	13,370	76	0.57	Riyadh	5,227	.	.
Los Angeles	13,223	30	0.23	Hangzhou	5,189	.	.
Rio de Janeiro	11,867	35	0.29	Dallas-Fort Worth	5,143	.	.
Manila	11,654	43	0.37	Singapore	5,086	78	1.53
Moscow	11,472	168	1.46	Chittagong	5,069	.	.
Osaka	11,430	125	1.09	Pune	4,951	.	.
Cairo	11,031	51	0.46	Atlanta	4,875	38	0.78
Istanbul	10,953	12	0.11	Xi'an, Shaanxi	4,846	.	.
Lagos	10,788	.	.	Saint Petersburg	4,842	63	1.30
Paris	10,516	299	2.84	Luanda	4,790	.	.
Guangzhou	10,486	123	1.17	Houston	4,785	.	.
Shenzhen	10,222	47	0.46	Boston	4,772	74	1.55
Seoul	9,751	360	3.69	Washington, D.C.	4,634	86	1.86
Chongqing	9,732	.	.	Khartoum	4,516	.	.
Jakarta	9,630	.	.	Sydney	4,479	.	.
Chicago	9,545	142	1.49	Guadalajara	4,442	17	0.38
Lima	8,950	16	0.18	Surat	4,438	.	.
London	8,923	267	2.99	Alexandria	4,400	.	.
Wuhan	8,904	25	0.28	Detroit	4,364	12	0.27
Tianjin	8,535	36	0.42	Yangon	4,356	.	.
Chennai	8,523	.	.	Abidjan	4,151	.	.
Bogota	8,502	.	.	Monterrey	4,100	32	0.78
Kinshasa	8,415	.	.	Ankara	4,074	12	0.29
Bangalore	8,275	.	.	Shantou	4,062	.	.
Bangkok	8,213	51	0.62	Salvador	3,947	.	.
Hyderabad	7,578	.	.	Melbourne	3,896	.	.
Lahore	7,352	.	.	Porto Alegre	3,892	17	0.44
Tehran	7,243	54	0.75	Phoenix	3,830	.	.
Dongguan	7,160	.	.	Montreal	3,808	68	1.79
Hong Kong	7,053	54	0.77	Zhengzhou	3,796	.	.
Madrid	6,405	239	3.73	Johannesburg	3,763	.	.
Chengdu	6,397	16	0.25	Brasilia	3,701	27	0.73
Ahmadabad	6,210	.	.	Recife	3,684	28	0.76
Foshan	6,208	.	.	San Francisco	3,681	48	1.30

Note: Populations in thousands. Subway stations per person is per 100,000 residents.

that are topcoded in DMSP-OLS version. Fewer cross-section of the radiance calibrated lights are available, but fortunately, the available cross-sections are about 1995, 2000, 2005 and 2010, and so they match up neatly with the last four cross-sections of our population data.

To assign lights at night cells to cities, we must know the locations of cities. In addition to population data, the UN World Cities Data provide the latitudes and longitudes of city centers. Inspection of these centers against city centers identified by other means indicates that, with a few exceptions that we correct, they are quite accurate.[8] Given this, we also use the UN World Cities Data to locate our cities. Figure 1 illustrates the 2010 radiance calibrated lights at night data for six cities. The figure also shows 5km and 25km circles centered on city locations given by the UN World Cities Data, and the city's subway network. In every case except for Mexico City, location of the center recorded in the UN World Cities Data corresponds to the center of the lights data and the subway network. For Mexico City, the recorded in the UN World Cities Data is a few kilometers too far north. While we occasionally measure the location of city centers with some error, overall our data seems quite accurate.

Lights at night data are of interest as a check on our population data. The lights at night data are measured consistently across cities and we can calculate city level measures of total light without reference to administrative boundaries. That is, the lights at night data are not subject to either of the two problems that we are concerned about for our population data. Since people light the places they live and work, more densely populated and more productive places are often brighter. More concretely, Henderson and Storeygard 2012 use the topcoded version of lights at night data to show that country level mean light intensities are a good proxy for GDP, a result that Storeygard 2012 confirms at the regional level for China.

The second to last line of each panel of table 1 shows the correlation of the mean light intensity within 25km of a city center in 2010 and population in 2010 in the whole sample and in the subsample of subway cities. It is clear that lights provide some information about population, although this information is imperfect. Finally, we note that the lights at night data are difficult to interpret. While we can be confident that lights at night data are telling us something about the location of economic activity, we cannot know whether places are brighter because the people living there are richer, because the place is more densely populated, or because it is the site of a large office tower or factory.

Perhaps more importantly, the fine spatial resolution of the lights at night data allows us to examine the spatial structure of cities. In particular, we are able calculate measures which indicate the extent to which activity in the city is centralized or diffuse. For example, increases in the share of all light within 25km of the center that is within 5km indicates a city where activity is more

[8] We checked the reasonableness of the UN cities locations in two ways. First, we hand checked several against other sources of data, google maps and lights at night. Second, for China we have a competing set of centers based on the brightest point in the lights at night data (see Baum-Snow et al 2012 for details), and these centers matched the UN centers closely.

concentrated in the downtown core. By using these sorts of measures, we are able to investigate the relationship between subways and urban form.

4. The relationship between subways and population

We now turn to a description of the relationship between subways and population. Figure 4 shows the relationship in 2010 between city size and the incidence of subway systems for all of the cities in our sample excluding Tokyo.[9] The horizontal axis gives city population by 0.5m bin and the vertical axis gives the proportion of cities with subways for each bin. We split our sample of cities into rich and poor country cities on the basis of the IMF advanced economy list for 2012.[10] Grey squares and black triangles indicate the share of rich and poor country cities with subways. The markers are spaced irregularly along the horizontal axis because some population bins are empty. The solid line is a smoothed plot of subway frequency in rich country cities and the dashed line is the corresponding plot for poor country cities.[11]

There are no rich country cities with population above 5m without a subway system and subways are common even among rich country cities with populations in the 1m-5m range. Subways are relatively rare among developing cities with populations less than about 5m and their frequency increases more or less smoothly with city size.

Table 2 describes the largest 90 cities in our sample as of 2010. For each city, the table reports population, the count of operational stations and the number of stations per 100,000 of population. Unsurprisingly, world destination cities are well represented among subway cities. Paris, London, New York, Washington D.C. and Shanghai all have subways. On the other hand, and perhaps more interestingly, many of the subway cities in table 2 do not often attract attention on the world stage or appear in travel advertisements. The notion that building a subway will cause a city to join the ranks of 'world cities' seems inconsistent with this table.

Moreover, despite the strong relationship between city size and the presence of a subway system that we see in figure 4, table 2 suggests that the relationship between population and subways is nuanced. In particular, none of the four cities larger than New York has even half as many subway stations and Tokyo has fewer than one third as many stations per person as New York. Looking down the list, we see that such reversals are common and do not simply reflect rich and poor country differences. Consistent with this, the raw correlation between operational stations and population in 2010 is about 0.58. While subways are clearly more common in big cities, the relationship between system size and city size is noisy. Table 2 further suggests that subway

[9] At 36 million people, Tokyo is nearly twice as large as the second largest city. We omit it from the figure to improve legibility.

[10] These rich countries are: Australia, Japan, New Zealand, the United States, Canada, Austria, Belgium, Czech Republic, Denmark, Finland, France, Germany, Greece, Ireland, Israel, Italy, Netherlands, Norway, Portugal, Singapore, South Korea, Spain, Sweden, Switzerland and the United Kingdom.

[11] More specifically, both lines are kernel weighted local polynomial regressions.

Table 3: Mean city-year population growth rates by time to a subway expansion

	$t-2$	$t-1$	t	$t+1$	$t+2$	N
Panel a						
			0.063	0.054***		138
			0.078	0.067**	0.064**	60
		0.090***	0.073			204
	0.120**	0.107***	0.083			141
		0.075***	0.061	0.052**		64
Panel b						
				−0.001		138
				−0.001	−0.009	60
		0.006*				204
	0.013*	0.012*				141
		0.009*		−0.007		64

Notes: Each row in panel (a) shows growth rates of cities in consecutive time periods. Each row in panel (b) shows the difference in growth rates of cities (relative to period t) in consecutive time periods from a regression controlling for year×continent dummies. t is a period of subway expansion. $j \neq t$ is a period with no subway expansion. Stars indicate a significant difference of growth rate compared to period t. *** 1%, ** 5%, * 10% significance respectively.

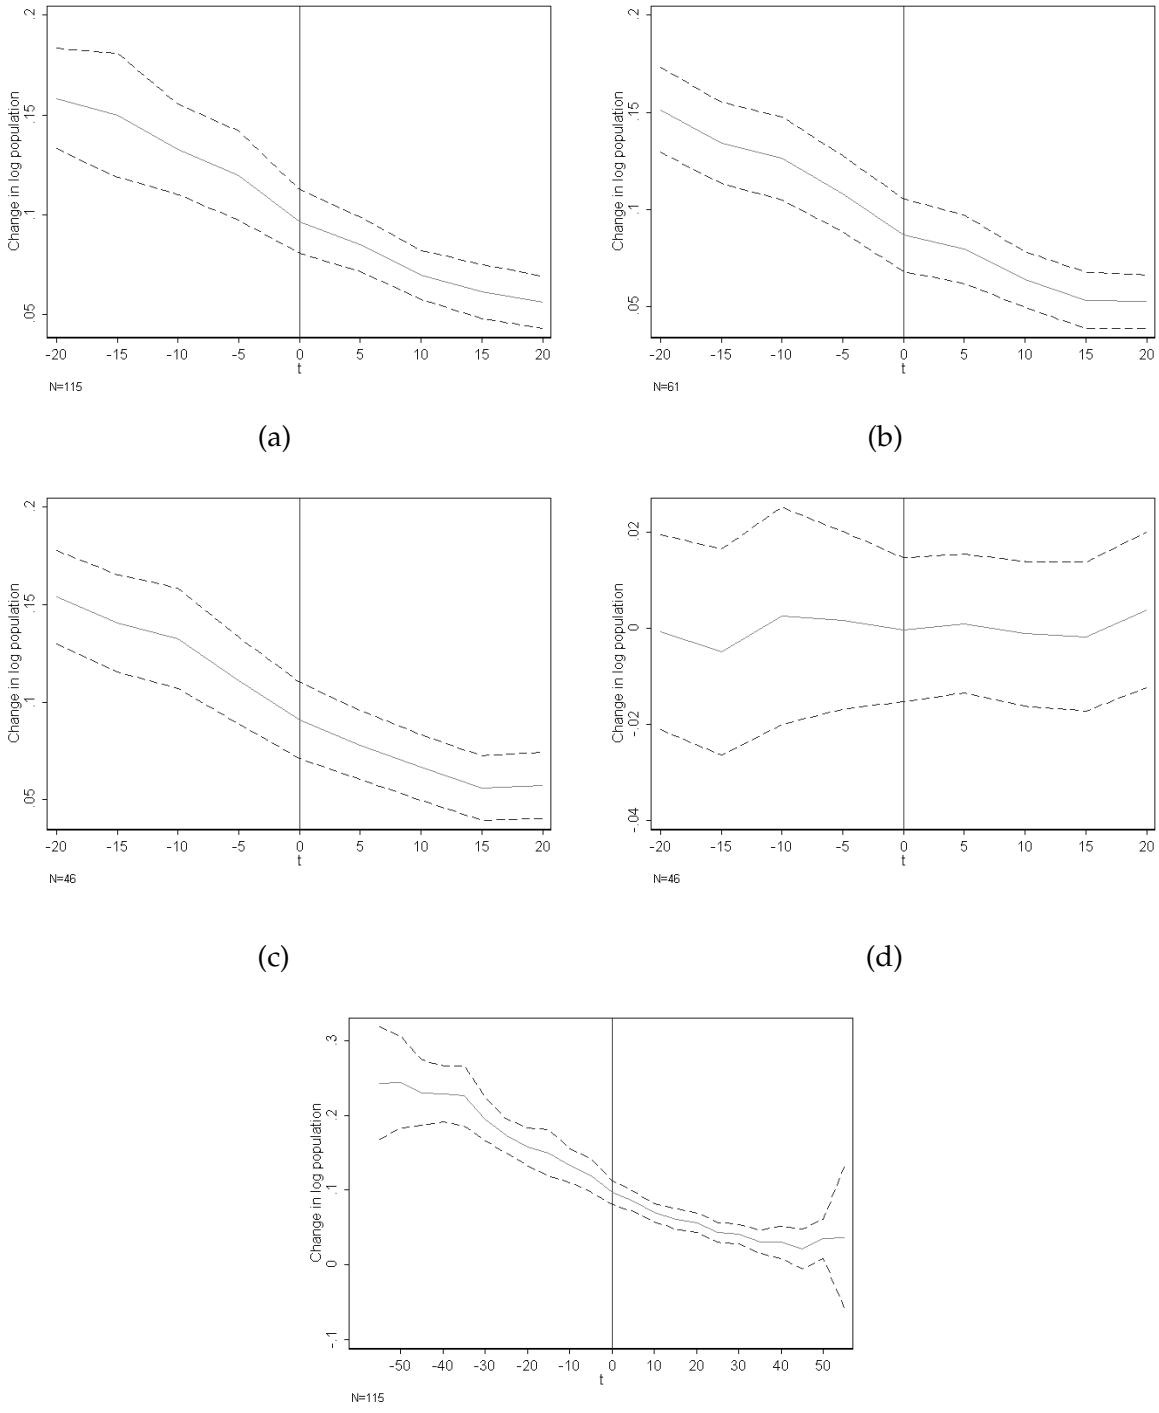
capacity is rarely, if ever, a binding constraint on city size. Indeed, some of the world's largest cities have no subway system to speak of.

We now turn to an investigation of what happens to a city when its subway system changes. Figure 5 presents five panels describing the relationship between changes in population and subway system extent. The horizontal axis of panel (a) indicates time in years since a subway system opening, with negative values indicating years prior and positive values years after. The vertical axis indicates the mean change in log population, i.e., the population growth rate, for all cities during the five year period ending t years before or after the subway opening. The solid line plots the mean growth rate and dashed lines give upper and lower 95% confidence bounds.[12] From figure 5 panel (a), the average population growth rate of cities during the five years following the opening of their subway systems is about 8%. During the five year period preceding a subway opening by five years, the average population growth rate is about 12%. Panel (a) of figure 5 is unambiguously declining. During the 20 years before and after a subway opening, the average city in our sample sees its growth rate decline and there is no obvious change in this trend around the opening of the subway system.

Panel (a) is based on all 115 cities in our sample where subway systems begin to operate between 1930 and 2010. Since many cities in our sample open their subway systems before 1970 or after 1990 we cannot calculate their growth rates both 20 years before and after their system opening. Thus,

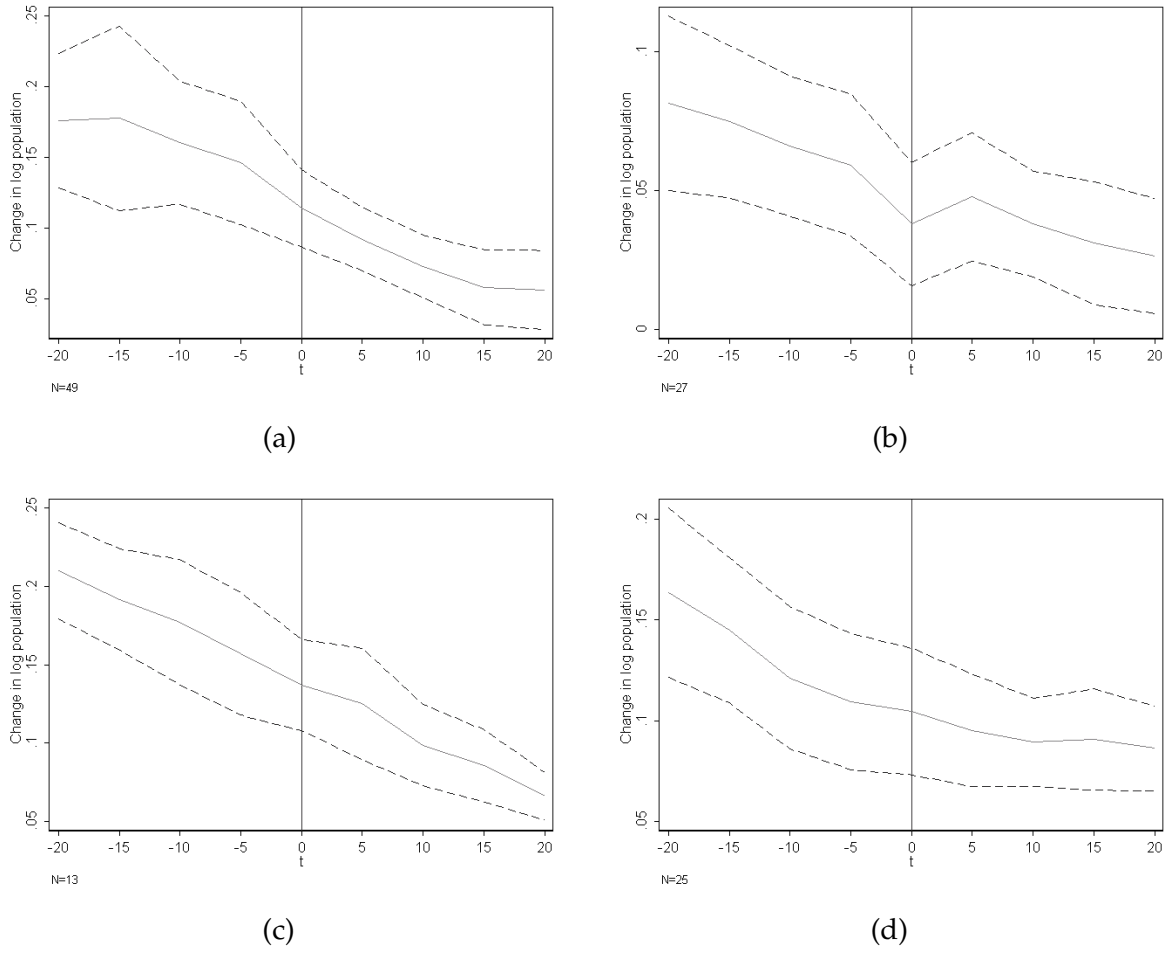
[12] These are local bounds constructed by connecting upper and lower 5% bounds at each year.

Figure 5: Subways openings and expansions, and growth in population and lights



Note: (a) Mean population growth rates by time from system opening, all subway cities. (b) Mean population growth rates by time from system opening, constant sample. (c) Mean population growth rates by time from system opening, constant sample of cities with pop >1m in 1970. (d) Mean deviation from annual population growth rates by time from system opening, constant sample of cities with pop >1m in 1970. (e) Mean population growth rates by time from system opening, long time horizon.

Figure 6: Subways openings and growth in population by continent



Note: Mean population growth rates by time from system opening, all subway cities in (a) Asia (b) Europe (c) North America (d) South America.

the growth rates for the different times in panel (a) do not reflect the same set of cities. In addition, the UN World Cities Data includes all cities with a population over 750,000 at any time during 1950-2010. This means that our sample over-represents fast growing cities that were small in the early part of the sample, but grew rapidly during the study period.

Panels (b) and (c) of figure 5 investigate the importance of these sampling problems. Panel (b) replicates panel (a) using the constant sample of cities for which we can calculate 20 years of population growth before and after their subway system opening, i.e., cities whose subways opened between 1970 and 1990. Panel (c) further restricts attention to cities with populations over 1m in 1970. This eliminates the small fast growing cities that qualify for the sample late in the sampling period. Panels (b) and (c) show the same pattern that we saw in panel (a).

The downward trend described in panels (a), (b) and (c) probably reflects the sample wide decrease in growth rates documented in table 1. It may be that this downward trend masks increases in growth rates associated with subway system openings. Panel (d) of figure 5 investigates this possibility by controlling for each period's mean growth rate. For this figure, for each year we calculate each city's residual growth rate from a regression of growth rates on continent specific year dummies. We next calculate the average of these residuals as a function of time from subway opening. To avoid sampling problems, we continue to restrict attention to the sample from panel (c), cities over 1m in 1970 with subway openings between 1970 and 1990. Unsurprisingly, this process removes the downward trend that we see in the first three panels, but still does not show a systematic change in growth rates following subway system openings.

The first four panels of figure 5 show that city population growth rates do not increase during the 20 year period following the opening of a subway system. As we discuss in section 2, the literature documents effects of subways on within city outcomes over much shorter periods and the effects of other types of infrastructure on city level outcomes over a 10-20 year horizon. Thus, the 40 year period illustrated in figure 5 (a)-(d) should be long enough to reveal whether growth rates respond to a subway system opening. Nevertheless, in panel (e) we use our entire sample of cities and investigate population growth rates over the longest time period that our 60 year sample allows, 55 years. This figure suggests that the pattern we see in panel (a) extends nearly sixty years before and after a subway opening, although our estimates become noisier as the time from the subway opening approaches 55 years.

Finally, to check for differences across regions in the relationship between urban growth and subways, figure 6 reproduces panel (a) of figure 5 for subsets of subway cities in Asia, Europe, North America and South America.[13] Remarkably, each of the continents shows a similar pattern. Urban population growth rates are declining in the period around subway openings and there is no obvious sign of a change in this trend at the time a subway opens. The only qualification of this statement applies to Europe, where there is a statistically insignificant deviation from trend around the opening of a subway system.

[13] We exclude Australia and Africa because they have zero and one subway city, respectively.

Figure 5 describes population growth rates as time varies relative to the date of a subway system *opening*. In Table 3 we turn our attention to the relationship between subway *expansions* and growth rates. The top row of panel (a) describes 138 city-year pairs where a city-year with a subway expansion is followed by a city-year without a subway expansion. On average, the growth rate in city-years with an expansion is 0.063, and in the subsequent city-year, without an expansion, it is 0.054. A t-test of the difference between the two means indicates that they are statistically different with high probability. In short, population growth rates are lower following a subway expansion than during one.

The remaining three rows of panel (a) of table 3 perform similar calculations for slightly different sets of city-years. In row two we consider the 60 city-year triples for which we observe a subway expansion followed by two city-years without an expansion. As for row 1, we see that growth rates decline following a subway expansion and that the decrease in growth rate is statistically different from zero. In the third row we consider the 204 pairs of city-years where a subway expansion follows a city-year without an expansion. The mean growth rate for city-years preceding a subway expansion is larger than for city-years with an expansion, and this difference is statistically different from zero. The fourth row of table 3 considers the 141 triples of city-years where a subway expansion is preceded by two years without an expansion. Again, we see that city growth rates decline in the years leading up to a subway expansion. The last row of table 3 considers the 64 triples of city-years for which a subway expansion follows and precedes city-years without expansions. The pattern of the other rows is preserved. Population growth rates are higher before a subway expansion and lower after, and this trend is statistically different from zero.

Similarly to panel (d) of figure 5, panel (b) of table 3 replicates the results of panel (a), but controls for continent specific year effects. Specifically, the values reported in panel (b) of table 3 are regression coefficients β from the regression,

$$\Delta log(\text{Pop}_{it}) = \alpha_{jt} \cdot \text{Continent}_j \times \text{Year Indicators}_t + \beta \cdot \text{Time to Expansion Indicators}_{it} + \epsilon_{it},$$

which we estimate with robust errors, clustered at the city level, using the same samples as in the top panel. We test whether the various time to expansion coefficients are different from the year zero coefficient using a robust F-test. Even after we control for the continent specific year effects, subway expansions are not associated with a measurable change in population growth rates.

5. Econometric model

The descriptive evidence presented so far indicates a positive cross-sectional relationship between the extent of a city's subway network and its population. Larger cities have more extensive subway networks. On the other hand, time series evidence suggests that incremental changes to subway networks do not affect the population of cities. The simplest explanation of these two findings is that large cities build and expand subway networks but that these networks do not

affect subsequent population growth. Our goal in this section is to subject this explanation to more exhaustive testing. We are also interested in exploring the dynamics of how city population levels respond to subways more carefully.

To proceed, index the set of observed cities by $i \in I$ and the set of observed years by $t \in T$. Let y_{it} denote an outcome of interest for city i in year t. Depending on context, y will be population, light intensity within a radius of the city center, or a measure of the centrality of the city based on lights data. Let s_{it} denote a measure of subway extent in city i in year t, usually the number of operational stations but sometimes the number of operational subway lines or route kilometers. Let x_{it} denote a vector of time varying city level covariates, most often country level GDP per capita and continent specific year indicators, and z_i a time-invariant vector of city level controls describing physical geography. To allow the interpretation of regression coefficients as elasticities, quantities are typically in logarithms and where necessary we add one to variables to facilitate this transformation. We use the operator Δ to denote first differences, so that $\Delta x_t = x_t - x_{t-1}$.

We are concerned that equilibrium subway construction occurs in slow growing cities or city-years. In this case, non-randomness in the assignment of subways to cities may mask a positive causal relationship between subway extent and urban population growth. We are specifically concerned with two possibilities. First, that subway construction is correlated with time-invariant unobserved factors that affect the city population. We can address this possibility by including city specific intercepts in our regressions in the standard way. Second, we are concerned that subway construction systematically occurs in response to a sequence of improbably large population increases, so that a decrease in population growth rates naturally coincides with subway expansions. Alternatively, if subway construction draws people to the city who leave when construction is complete, we will tend to observe subway expansions in years with negative population shocks. This second inference problem arises because subway construction is caused by a shock or a series of shocks to recent population. To resolve this problem, we examine the effects of subways predicted by lags of subways that predate the hypothetical sequence of problematic shocks.

The following system, while too stark to be defensible, formalizes this intuition and allows us to illustrate the econometric problems that arise in its implementation,

$$y_{it} = A_1 s_{it} + c_i + \epsilon_{it} \quad (1)$$

$$s_{it} = B_1 s_{it-k} + d_i + \eta_{it}, \quad (2)$$

where A_1, the "outcome elasticity of subway extent", is the parameter of interest and k is a positive integer. In words, population depends on contemporaneous subways, a city specific intercept and a random disturbance. Subways at t depend on subways at period $t-k$, a city specific intercept and a random disturbance.

Written this way, it is natural to consider using s_{it-k} as an instrument for s_{it}. This is subject to two objections. First, this system of equation commits us to a particular dynamic structure for the

relationship between subways and population. It is natural to wonder whether this dynamic structure is correct. In our estimations we consider alternative dynamic structures for our data. Second, evidence presented so far suggests that the unobserved determinants of subway construction are related to unobserved determinants of growth. That is, $cov(c_i, d_i) \neq 0$. It follows that, because s_{it-k} also depends on d_i, we should not expect $cov((c_i + \epsilon_{it}), s_{it-k}) = 0$. That is, the dynamic structure described by equations (1) and (2) requires that s_{it-k} be correlated with unobservables in the population equation, and thus, that it is not a valid instrument.

As a first response to this problem, first difference equations 1 and 2 to get

$$\Delta y_{it} = A_1 \Delta s_{it} + \Delta \epsilon_{it} \quad (3)$$

$$\Delta s_{it} = B_1 \Delta s_{it-k} + \Delta \eta_{it}. \quad (4)$$

Differencing solves two problems. First, and as usual, it removes time-invariant unobservables from the first equation. Second, after removing the city specific intercept from the population equation, the validity of lagged subways as an instrument for current subways hinges on the whether $cov(\Delta s_{it-k}, \Delta \epsilon_{it}) = 0$, or in words, on whether lagged change in subways is uncorrelated with current change in the time varying propensity to grow. This is simply a more technical statement of the intuition that motivates this instrumental variables strategy. We are concerned that subway construction reflects a series of recent population shocks. To resolve this problem, we use subways predating the hypothetical sequence of problematic shocks to predict subways at time t. This strategy is valid if the shocks that determine lagged subways are uncorrelated with the shocks that determine current population growth.[14] In sum, this discussion suggests that we can use long lags of subways as instruments for current subways if we implement a rich enough dynamic model and if we difference to eliminate city specific effects. Note that, as is standard in panel data models, in place of lagged changes of subways, we can (and do) use lags of the component levels as our instruments without affecting the basic intuition of this estimation strategy.

On the basis of the discussion above, except for special cases, a necessary condition for $E(\Delta s_{it-k} \Delta \epsilon_{it}) = 0$ is that

$$E(\Delta \eta_{it-k} \Delta \epsilon_{it}) = 0. \quad (5)$$

That is, no serial correlation of lagged differenced errors in the subways equation with current differenced errors in the population equation. Such serial correlation of errors is easily testable and on the basis of the intuition above, we regard it as a specification test in instrumental variables estimations. In many of our estimations we report tests for serial correlation of ϵ_{it}, which we also regard as specification tests. The rationale for this is heuristic; the presence of serial correlation

[14] While differencing solves one problem, it may create another. If $k = 1$ then both Δs_{it-1} and Δy_{it} involve terms for quantities for time t. If we are concerned about contemporaneous correlation of errors in the population and subway equations, then this creates an obvious problem. This is a classic problem in dynamic panel data and the conventional approach is to substitute s_{it-2} for Δs_{it-1} or to use longer lags.

indicates that we have not completely described the dynamics of the relationship between population and subways.[15]

The discussion above describes an econometric strategy based around using old subways to instrument for current subways. An alternative is to use old levels or changes of population to instrument for current changes in subways. The basic logic of this approach is similar to that described above. However, lagged population levels and changes have less ability to predict current changes to subways than do lagged subway variables, so we organize our discussion and analysis around the lagged subways instruments.

The instrumental variable strategy articulated above responds to the possibility that subway construction reflects recent trends in population. A similar problem arises if both population growth and subway growth reflect some unobserved city specific factor. For example, it may be that badly governed cities grow slowly and build subway networks. In this case, our estimated effect of subways on population growth confound the effects of bad municipal government with the effects of subways. To address this possibility, we would like to include fixed effects in the first differences regressions, or equivalently, city specific trend in the levels regressions, equations (1) and (2). To implement this estimator, we second difference equation (1).[16]

Summarizing, our econometric investigation will be organized around estimating the following system,

$$y_{it} = A_1 s_{it} + A_2 x_{it} + A_3 z_i + c_i + g_i t + \mu_{it} \quad (6)$$

$$s_{it} = B_1 s_{it-k} + B_2 x_{it} + B_3 z_i + d_i + h_i t + \eta_{it}. \quad (7)$$

This generalizes equations 1 and 2 in a number of ways. First, it allows for time-invariant control variables, z_i. Second, it allows for city specific trends and intercepts in both population and subways equations. Third, it allows for time varying controls. In practice, we will usually predict current changes in subways with 20 year old subways changes or levels, so that $k = 4$.

We are also concerned with two other inference problems. These are sampling problems and omitted variables bias. First, by construction, time series evidence on the effect of subways relies exclusively on cities that build subways, while cross-sectional evidence may include cities without subways. Second, our sample includes cities whose population exceeds 750,000 at any time during 1950-2010. As we discussed above, this leads to an oversampling of small, fast growing cities. To address these problems we experiment with a number of different sampling rules.

It is also possible that cities that build subways systematically engage in some other policy that affects growth. For example, if subway system construction was always accompanied by dramatic tax increases, then we would estimate the joint effect of subways and tax increases,

[15]We note that the instrumental variables strategy described here is related to the one proposed by Olley and Pakes (1991), while the exogeneity condition of equation 5 is related to ideas developed in Arellano and Bond (1991).

[16]In principle, one could also implement our instrumental variables strategy in second differences. We experimented with this but found that lagged subways and population variables do not have much ability to predict current second differences of subways. Consequently, these regressions were not informative.

not simply the effect of subways, as we intend. Typical responses to omitted variables problems are essentially the same as those to the endogeneity problems that we discuss above. That is, fixed effects and instrumental variables estimation. With this said, if we regard our differencing and instrumental variables strategies as a response to omitted variables, then the rational for this strategy is somewhat different. In particular, instrumenting for current changes in subways with old subways isolates the effects of subways from any local public policy with different timing than subway construction.

6. Subways and population: Main estimation results

We proceed by estimating successively more complete and complex versions of equations 6 and 7. To begin, in table 4 we estimate equation 6 using OLS on pooled cross-sections. Such estimations result in unbiased estimates only if the error terms are uncorrelated with the explanatory variables. That is, if $E(\epsilon_{it} + c_i + g_i t | x_{it}, z_i, s_{it}) = 0$. This condition seems implausibly strong. We expect that unobserved factors affecting the productivity or attractiveness of a city also affect its subways, so we regard these estimations as primarily descriptive.

In column 1 of table 4 we regress the log of population on log of the count of operational subway stations. We use the entire sample of 632 cities for which we have population and subway data. Since our panel is complete for these two variables, we have a sample of $13 \times 632 = 8216$ city-years. The subway elasticity of population is large. A 10% increase in a city's count of stations is associated with a 4.8% increase in population. Column 2 replicates this result, but controls for country level GDP and continent-by-year fixed effects, along with several time-invariant controls; a capital city indicator, and distances to the ocean, international boundary and nearest navigable river. We see that the coefficient on subways, while still large, decreases to 0.31. Our sample size decreases to 7396 in this regression, primarily because a number of the countries covered by our sample, particularly those in the former Soviet Union, came into existence after 1950 and so country level GDP is not available.

Column 3 considers the same regression as column 2, but restricts attention to cities that had subways in 2010. This is the largest sample of cities that could possibly contribute to a first differences estimate of the effect of subways. This reduces our sample size to 1565 city-years, but leaves the coefficient of subways unchanged. The sample of 137 cities used in column 3 includes some cities that were small in 1950 and grew quickly to cross the 750,000 threshold for inclusion in the UN World Cities Data. Column 4 restricts attention to cities which had a subway in 2010 and a population of more than 1m in 1970. This will be our preferred sample since it does not oversample fast growing cities and contains only cities where we observe changes in subway extent. Unsurprisingly, the subway coefficient is a bit smaller in this sample. Columns 5 and 6 replicate column 4, but use different measures of subways as the explanatory variable. Columns 5 and 6 consider alternative measures of subway extent, route kilometers and log subway

lines. Coefficient magnitudes change approximately in proportion to the changes in the standard deviation of the subway measures. Thus, our cross-sectional estimates are not artifacts of our particular measure of subway extent.

In column 7 we restrict attention to city-years where a subway has been in operation for at least twenty years. This reduces our sample to the 75 cities where subways operated during or prior to 1990 and to the 355 city-years that post-date subway opening by twenty or more years. This sample is of interest both because it suggests that old subways may be somewhat more important than new, and because it will help us to interpret instrumental variables estimates presented later. Finally, column 8 reports a regression similar to column 4, where our dependent variable is the logarithm of mean light intensity in a 25 km disk centered on the city. As in column 4, we restrict attention to cities with subways in 2010 and a population above 1m in 1970. Our sample of city-years is smaller than for population regressions, even though the sample of cities is the same. This is because we have just four cross sections of lights data. We see that a one percent increase in subways causes a 0.22 percent increase in lights. This is almost identical to our results for population and suggests that our population regressions are not driven by problems in the UN World Cities Data.

In sum, table 4 shows what we expect on the basis of figure 4. Cities with more subways tend to be bigger. This relationship is robust to time-invariant controls, sampling, the particular measure of subway extent and whether we measure city size with lights or population. In all of the regressions reported in table 4 we test for and fail to reject serial correlation of order 1, and of orders 2 and 3 for the longer population panels. This indicates the presence of a dynamic structure not described by the pooled cross-section specification.

We now turn to the problem of unobserved characteristics. Table 5 presents first difference estimates of a version of equation 6 without city specific trends. We note that both first difference and within estimators are consistent estimators for equation (6) if the errors, ϵ_{it} in each period are not correlated with the regressors in any period conditional on the unobserved effect. Because our approach to estimating equations 6 and 7 revolves around first difference estimations, we prefer the first differences estimator.[17]

Columns 2-4 and 8 in table 5 use the same sample as column 4 of table 4, while column 1 uses the slightly larger sample available when we do not control for changes in GDP. That is, we restrict attention to cities with subways in 2010 and with a population above 1m in 1970. In column 1, we report the results of regressing change in log population on change in the log of the count of operational stations. In column 2 we repeat this regression with continent specific year dummies

[17]The choice between the two estimators hinges on subtle differences in the errors. The first difference estimator is more efficient if ϵ_{it} is a random walk, while the within estimator is more efficient if the ϵ_{it} are i.i.d.. Apart from these two pure cases, serial correlation of the ϵ_{it} suggests the use of robust standard errors with either estimator. Second, and more substantively, if the data have a dynamic structure, then both first difference and within estimators are inconsistent (Ch. 10, Wooldridge (2001)).

Table 4: Pooled cross section

	All cities		Subway cities					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\ln(pop_t)$	$\ln(pop_t)$	$\ln(pop_t)$	$\ln(pop_t)$	$\ln(pop_t)$	$\ln(pop_t)$	$\ln(pop_t)$	$\ln(Lights_t)$
$\ln(s_t)$	0.48***	0.31***	0.30***	0.23***			0.54***	0.22***
	(0.02)	(0.03)	(0.03)	(0.04)			(0.07)	(0.04)
$\ln(\text{route km}_t)$					0.20***			
					(0.03)			
$\ln(\text{subway lines}_t)$						0.47***		
						(0.06)		
$\ln(GDPpc_t)$		0.18***	-0.18**	-0.21**	-0.20**	-0.22**	-0.29**	0.13
		(0.04)	(0.09)	(0.07)	(0.07)	(0.07)	(0.13)	(0.09)
Geographic controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
YearXContinent dummies	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep Variable	13.35	13.44	14.48	14.99	14.99	14.99	15.18	4.80
Mean of subways regressor	0.38	0.40	1.88	2.63	2.79	1.10	4.12	3.43
SD subways regressor	1.15	1.17	1.92	1.86	1.99	0.93	0.88	1.31
R-squared	0.18	0.48	0.46	0.49	0.47	0.52	0.67	0.51
Number of cities	632	629	137	99	99	99	75	99
Number of subway cities	138	137	137	99	99	99	75	99
Number of periods	13	13	13	9	9	9	9	4
p-value 1st order autocorr	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
p-value 2nd order autocorr	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
p-value 3rd order autocorr	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
Observations	8216	7396	1565	844	844	844	355	396

Dependent variable: Log population of metropolitan area in period t (except last column see (8) below).

City-level clustered standard errors in parentheses. Stars denote significance levels: * 0.10, ** 0.05, *** 0.01.

Geographic controls are continent dummy, capital city dummy, log km to ocean, log km to land border, and log km to major navigable river.

(1)-Pooled cross section. (2)-Add geographic controls, gdp pc control, and yearXcontinent dummies.

(3)-Only cities with subway by 2010. (4)-Cities with subway by 2010 and at least 1 million population in 1970, $t >=1970$.

(5)-Log route km as main regressor, $t >=1970$. (6)-Log subway lines as main regressor, $t >=1970$. (7)-Add $log(s_{t-4}+1) > 0$ restriction.

(8)-Dep. var. is log mean radiance calibrated lights in a 25km circle around the centroid of the city.

and change in country log GDP as controls. In both cases we estimate that doubling a city's stock of subways increases population by just less than 1%. These estimates are quite precise. In columns 3 and 4, we use route kilometers and lines as our measures of subway extent. The effect of a 1% increase in route kilometers is indistinguishable from the effect of a 1% increase in the count of stations. The effect of a 1% increase in the number of subway lines is somewhat larger than the effect of a 1% increase in the count of stations, but the standard deviation of subway lines is correspondingly smaller.

Column 5 is a long-differences specification. Here we use the same specification and sample of cities as in column 2, but take time differences between 1970 and 2010. This leads to a 40 year subway elasticity of population of about 9.5 percent. This is a much larger effect than we obtain on the basis of five year differences and could indicate a larger causal effect of subways on population. If we replicate the regression of column 5, but use change in subways between 1970 and 2000 as our explanatory variable, then the subway coefficient is 0.03 and is indistinguishable from zero. Thus, the larger subway coefficient in the long difference estimate seems to indicate that fast growing cities build subways rather than the converse.

In column 6 we use the average light intensity in a disk of 25km centered on the city as our dependent variable. As with our other regressions, we find a much smaller effect than in the comparable cross-sectional regression, column 8 of table 4, in this case not distinguishable from zero.

In column 7, we restrict attention to city-years where subways have been in operation for at least 20 years, the same sample as we used in column 7 of table 4. We see that the subways coefficient for this set of city-years is 0.056. This is not quite two standard deviations larger than the estimate for an average subway city. The primary interest of this specification is to help us interpret IV results, yet to be presented. With this said, this specification suggests the possibility that the effect of subways is not the same in all cities, an issue to which we return below.

Summing up, first difference estimates are dramatically smaller than cross-sectional estimates. For example, the first difference estimate from column 2 of table 5 is smaller by a factor of 30 than the cross-sectional estimate on the same sample in column 4 of table 4. If we use lights as our dependent variable the first difference estimates are smaller by a factor of 10. In total, these results suggest that unobserved city specific factors influence both population and subway extent.

In all of the regressions reported in table 4 we test for and reject no first-order serial correlation of orders 2 and 3 for the longer population panels. This indicates the presence of a dynamic structure not described by the first differences specification.[18] To investigate this feature of our data, and to deal with the endogeneity problem discussed in section 5, we now consider an IV regression that relies on long lags of subways as instruments. Table 11 in the appendix presents sample first stage results. Broadly, first stage results show that third and fourth lagged levels of subways are excellent predictors of current changes in subways, while lagged changes are somewhat worse, but

[18] Recall that first order serial correlation arises mechanically in first differences regressions.

still good. The results in table 11 are broadly consistent with figure 3. By inspection of this figure, the growth rate or size of a subway system twenty years ago allows a good guess at the current growth rate.

Columns 8 and 9 of table 5 report instrumental variables estimates of the effects of subways on population when we use the 4th lagged level of subways as our instrument for current changes. In column 8 we use the same sample as in column 2, all cities with subways in 2010 and population above 1m in 1970. The estimated effect of subways is about 0.07, almost 10 times as large as the corresponding estimate in column 2. In column 9, we restrict attention to city-years with subways for at least 20 years. Here we see that the estimated effect of subways is indistinguishable from zero.

These results require explanation. Our instrument is the fourth lag of subways. This is zero for any city-year in which the subway system is less than 20 years old. When we restrict attention to city-years in which predicted subways are non-zero in column 9, we see that there is no measurable effect of subways on population growth, although the first stage is strong. That is, the instrumented coefficient ceases to be significant once we focus on observations with positive predicted subways. Comparing columns 2 and 7 of table 5, we see that the effect of subways is larger for the sample of city-years with old subways than in the universe of city-years. This means that the large and significant IV coefficient in column 8 reflects differences between the sample with new subways (and hence no predicted subways and a small subways coefficient) and the sample with old subways. Thus, the relatively large estimated effect of subways in the IV regression reported in column 8 probably reflects this sampling issue rather than a causal effect of subways on population. We note that the instrumental variables estimates continue to show evidence of serial correlation, although this is less important for column 9 than for any other specification so far.

To perform the specification test suggested by equation 5, we calculate regression residuals from the first stage and structural equation associated with column 9 of table 5 and then regress structural equation errors on first stage errors. The p value of the regression coefficient is about 0.6. This allows us to reject the hypothesis that the two errors are correlated and provides evidence to support the exclusion restriction.

Table 6 presents second differences estimates of equation 6. Column 1 uses our preferred sample of all cities with subways in 2010 and population above 1m in 1970, and regresses the second difference of log population on the second difference of the log of operational stations. Column 2 controls for GDP and includes continent specific year indicators. In both regressions, the estimated effect of subways is indistinguishable from zero. Column 3 restricts attention to city-years where the subway system is at least 20 years old, the sample used in our preferred IV regression in column 9 of table 5. The estimated effect of subways on population growth is also zero in this

Table 5: First differences

	(1) $\Delta \ln(pop_t)$	(2) $\Delta \ln(pop_t)$	(3) $\Delta \ln(pop_t)$	(4) $\Delta \ln(pop_t)$	(5) $\Delta \ln(pop_{70-10})$	(6) $\Delta \ln(Lights_t)$	(7) $\Delta \ln(pop_t)$	(8) $\Delta \ln(pop_t)$	(9) $\Delta \ln(pop_t)$
$\Delta \ln(s_t)$	0.008**	0.007**				0.027	0.056**	0.074***	0.023
	(0.004)	(0.003)				(0.022)	(0.023)	(0.019)	(0.145)
$\Delta \ln(\text{route km}_t)$			0.007**						
			(0.003)						
$\Delta \ln(\text{subway lines}_t)$				0.019**					
				(0.009)					
$\Delta \ln(s_{70-10})$					0.095***				
					(0.025)				
$\Delta \ln(GDPpc_t)$		0.165***	0.164***	0.164***	0.806**	0.715***	0.098**	0.155***	0.107**
		(0.035)	(0.035)	(0.035)	(0.330)	(0.128)	(0.029)	(0.033)	(0.045)
YearXContinent dummies	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Geographic controls	No	No	No	No	Yes	No	No	No	No
Continent dummies	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep Variable	0.077	0.078	0.078	0.078	0.557	0.017	0.040	0.078	0.040
Mean of subways regressor	0.31	0.30	0.32	0.12	0.11	0.28	0.09	0.30	0.09
SD subways regressor	0.74	0.75	0.80	0.29	0.21	0.68	0.16	0.75	0.16
R-squared	0.01	0.40	0.40	0.40	0.56	0.52	0.39	0.08	0.38
Number of cities	99	99	99	99	99	99	75	99	75
Number of subway cities	99	99	99	99	99	99	75	99	75
Number of periods	9	9	9	9	1	3	9	9	9
p-value 2nd order autocorr	0.00	0.00	0.00	0.00			0.00	0.00	0.00
p-value 3rd order autocorr	0.00	0.00	0.00	0.00			0.05	0.01	0.10
F-stat excluded instrument								153.49	12.88
Observations	891	828	828	828	99	297	347	828	347

Dependent variable: Change in log population of metropolitan area in a 5 year period. Sample is large subway cities (over 1 million in 1970).

City-level clustered standard errors in parentheses. Stars denote significance levels: * 0.10, ** 0.05, *** 0.01. All regressions $t >= 1970$.

(1)- No controls. (2)-Add year dummies and change in log gdp controls.

(3)- Use change in log route km as regressor. (4)-Use change in log subway lines as regressor.

(5)- Long difference regression 1970-2010. (6)-Dep. var. is change in log mean radiance calibrated lights in a 25km circle around the centroid of the city.

(7)- Add $log(s_{t-4}+1) > 0$ restriction. (8) Instrument $\Delta \ln(s_t+1)$ with $log(s_{t-4}+1)$. (9) Add $log(s_{t-4}+1) > 0$ restriction.

Table 6: Second differences

	(1) $\Delta^2 \ln(\text{pop}_t)$	(2) $\Delta^2 \ln(\text{pop}_t)$	(3) $\Delta^2 \ln(\text{pop}_t)$	(4) $\Delta^2 \ln(\text{Lights}_t)$
$\Delta^2 \ln(s_t)$	-0.0021	-0.0014	-0.0077	0.0382
	(0.0014)	(0.0014)	(0.0093)	(0.0250)
$\Delta^2 \ln(\text{GDP}pc_t)$		0.0605***	0.0406**	1.0103**
		(0.0161)	(0.0197)	(0.3448)
YearXContinent dummies	No	Yes	Yes	Yes
Mean of Dep Variable	-0.009	-0.008	-0.003	-0.167
Mean of subways regressor	0.02	0.02	-0.01	-0.03
SD subways regressor	1.04	1.04	0.21	0.92
R-squared	0.002	0.184	0.314	0.407
Number of cities	99	99	75	99
Number of subway cities	99	99	75	99
Number of periods	9	9	9	2
p-value 3rd order autocorr	0.81	0.41	0.48	
Observations	891	812	338	198

Dependent variable: 2nd difference in log population of metropolitan area in a 5 year period.

Sample is large subway cities (over 1 million in 1970). All regressions $t >= 1970$.

City-level clustered standard errors in parentheses.

Stars denote significance levels: * 0.10, ** 0.05, *** 0.01.

(1)-No controls. (2)-Add year dummies and change in log gdp controls.

(3)-Add $log(s_{t-4}) > 0$ restriction.

(4)-Dep. var. is second change in log mean radiance calibrated lights in a 25km circle around the centroid of the city.

specification. Finally, column 4 uses log mean light intensity in a 25km disk centered on the city as the dependent variable. Again, the estimated effects of subways are indistinguishable from zero. First and second order autocorrelation arises mechanically in second differences estimations. The second difference regressions for population do not generally result in errors with 3st order autocorrelation. Thus, unlike our other estimations, the second difference regressions appear to provide a complete description of the dynamic structure of our data and suggest the presence of city specific unobservable characteristics correlated with population and subway growth.

We also experiment with regressions in levels rather than logs. These regressions generally lead to larger effects of subways. However, inspection of the data suggests that these effects are driven by a handful of observations. In fact, if we trim the top 5 percent of observations in our first difference regressions, the levels regression corresponding to column 2 of table 5 produces a subway coefficient not distinguishable from zero. The sample considered in column 7 of table 5 effectively restricts attention to larger cities and subways and yields a relatively large subways coefficient. Together with the results of our 'levels' regressions, this suggests that large cities may respond to subways differently than small cities.

Appendix table 12 investigates this possibility more carefully. In this table we partition the set of city-years used in column 7 of table 5, those for which our IV strategy is valid, into terciles on the basis of 1970 population and replicate our first difference and IV estimation for each subsample. These results support the idea that, in first differences, subways have a larger effect on big cities than small, though still less than 0.06. However, consistent with results in table 5 this coefficient is indistinguishable from zero in the corresponding IV regression. Here, however, we reach the limits of what our data can tell us about different size classes of subway cities. The first stage for IV results for the bottom two size terciles are weak, and so IV estimates for these size classes are not informative. To sum up, the data suggest that big cities may respond somewhat more positively to subways than smaller cities, even for big cities this effect is small, with elasticities on the order of 6 percent, and does not persist in our IV regression, which suggests that even this small effect is probably not causal.

Discussion We have presented four types of results, cross-sectional, first difference, IV and second difference. Consistent with descriptive evidence presented in section 1, cross-sectional estimates are much larger than first differences estimates. Results based on metropolitan area light intensity are qualitatively similar to those based on population. In our preferred sample, the cross-sectional estimate of the effect of doubling subway stations is a 23% increase in population. In first differences, the corresponding estimate is less than 1%. Our preferred IV result gives an estimated elasticity indistinguishable from zero as do our second difference estimates. All estimations except for those in second differences result in serially correlated errors. This favors the second difference estimates, while precision favors first differences. There is weak suggestive evidence that subways have larger, although still small effects, in larger cities.

Broadly, formal econometric results support the conclusion suggested by the descriptive evidence. That is, that big cities build subways and that these subways subsequently have little or no effect on the population in these cities. Our most favorable first difference and IV logarithmic regressions indicate that doubling a subway system will increase population by 5-7%. On the other hand, first difference estimates based on the whole sample of cities with subways in 2010 and population above 1m, our preferred IV estimate and all second difference estimates indicate that doubling subways causes a less than 1% increase in city population.

7. Subways and population: Extensions

We now consider three extensions to our main results. First, we investigate the possibility that subways affect cities in different regions differently. We present these results in table 7. In this table we replicate the regression presented in column 2 of table 5 for different subsets of cities. In column 1 we regress the first difference of log population on the first difference of the log of the count of operational stations, along with the first difference in log country GDP and continent specific year dummies. We restrict attention to the 38 cities in Asia with subways in 2010 and population above 1m in 1970. Columns 2,3 and 4 repeat this exercise for the corresponding cities in Europe, the US and Canada, and Latin America. The coefficient on subways ranges between 0.004 and 0.008, close to what we obtain for the whole sample, although standard errors are somewhat larger. This does not suggest that subways have dramatically different effects across different regions. In columns 5 and 6, we divide cities into those that lie in poor countries, again using the IMF definition from figure 4 and the complementary set in rich countries. Here we see that the effect of subways on population growth in rich countries is indistinguishable from zero and is estimated precisely: a two standard deviation increase in the subways coefficient still gives an elasticity of less than 1%. The effect of subways on poor country city population seems to be marginally larger, with an elasticity of just over 1%. Although this is still a small effect, that the effect of subways should be larger in poor countries seems consistent with the fact that subways are more heavily used in poor neighborhoods in rich countries (Gordon and Willson 1984; Glaeser, Khan, and Rappaport 2008).

We next consider models that allow for a richer dynamic structure in our data. Table 8 presents a series of estimates based on our preferred specification in column 2 of table 5, but using different lags of the subways variable. In column 1, we replicate column 2 of table 5 and in columns 2-4 we substitute successively older lags of subways. Interestingly, while the current change in subways has a small effect on population, older changes in subways do not. In columns 5 and 6 we include more than one lag of subways and see that coefficients are virtually identical to those we obtained when we include subway variables one at a time. This suggests that city populations adjust to new

Table 7: First Differences - Heterogenous treatment effects

	Asia	Europe	US & Canada	Latin America	Rich country	Poor country
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta \ln(\text{pop}_t)$	$\Delta \ln(\text{pop}_t)$	$\Delta \ln(\text{pop}_t)$	$\Delta \ln(\text{pop}_t)$	$\Delta \ln(\text{pop}_t)$	$\Delta \ln(\text{pop}_t)$
$\Delta \ln(s_t)$	0.007	0.005	0.004	0.008**	0.002	0.011*
	(0.006)	(0.007)	(0.004)	(0.004)	(0.003)	(0.006)
$\Delta \ln(\text{GDP}pc_t)$	0.171***	0.156**	0.019	0.117*	0.178**	0.156***
	(0.041)	(0.052)	(0.085)	(0.061)	(0.052)	(0.044)
YearXContinent dummies	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep Variable	0.11	0.02	0.07	0.09	0.05	0.11
Number of cities	38	27	20	33	54	45
Number of subway cities	38	27	20	33	54	45
Number of periods	9	9	9	9	9	9
Observations	301	221	180	297	478	350

Dependent variable: Change in log population in a 5 year period.

Sample is large subway cities (over 1 million in 1970). City-level clustered standard errors in parentheses.

Stars denote significance levels: * 0.10, ** 0.05, *** 0.01. All regressions $t >= 1970$.

(1) Asia (2) Europe (3) US & Canada (4) Latin America (5) IMF high income countries (6) Low income countries.

Table 8: First Differences - Distributed Lag Models

	(1) $\Delta \ln(pop_t)$	(2) $\Delta \ln(pop_t)$	(3) $\Delta \ln(pop_t)$	(4) $\Delta \ln(pop_t)$	(5) $\Delta \ln(pop_t)$	(6) $\Delta \ln(pop_t)$	(7) $\Delta \ln(pop_t)$
$\Delta \ln(s_{t+3})$							0.013**
							(0.005)
$\Delta \ln(s_{t+2})$							0.012**
							(0.005)
$\Delta \ln(s_{t+1})$							0.007
							(0.005)
$\Delta \ln(s_t)$	0.007**				0.007**		0.007*
	(0.003)				(0.003)		(0.004)
$\Delta \ln(s_{t-1})$		0.004			0.004		0.005
		(0.003)			(0.003)		(0.004)
$\Delta \ln(s_{t-2})$			-0.000		0.000	-0.000	0.004
			(0.003)		(0.003)	(0.003)	(0.004)
$\Delta \ln(s_{t-3})$				-0.003	-0.002	-0.003	0.003
				(0.003)	(0.003)	(0.003)	(0.004)
$\Delta \ln(\text{GDP}pc_t)$	0.165***	0.165***	0.166***	0.166***	0.164***	0.165***	0.129***
	(0.035)	(0.036)	(0.036)	(0.036)	(0.035)	(0.036)	(0.031)
YearXContinent dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep Variable	0.08	0.08	0.08	0.08	0.08	0.08	0.09
Number of cities	99	99	99	99	99	99	96
Number of subway cities	99	99	99	99	99	99	96
Number of periods	9	9	9	9	9	9	6
Observations	828	828	828	828	828	828	531

Dependent variable: Change in log population in a 5 year period.
Sample is large cities (over 1 million in 1970) with subway in 2010. City-level clustered standard errors in parentheses.
Stars denote significance levels: * 0.10, ** 0.05, *** 0.01. All regressions $t >= 1970$.

Table 9: First Differences - Regional population

	Pooled OLS		FD Estimation			
	(1)	(2)	(3)	(4)	(5)	(6)
	$\ln(\text{region pop}_t)$	$\ln(\text{region pop}_t)$	$\Delta\ln(\text{region pop}_t)$	$\Delta\ln(\text{region pop}_t)$	$\Delta\ln(\text{region pop}_t)$	$\Delta\ln(\text{region pop}_t)$
$\Delta\ln(s_t)$			0.012**	0.009**	0.006*	
			(0.004)	(0.004)	(0.003)	
$\Delta\ln(pop_t)$					0.455***	0.460***
					(0.070)	(0.070)
$\Delta\ln(\text{GDP}pc_t)$				0.246***	0.171***	0.171***
				(0.041)	(0.044)	(0.044)
$\ln(s_t)$	0.254					
	(0.188)					
$\ln(\text{route km}_t)$		0.231				
		(0.183)				
$\ln(\text{GDP}pc_t)$	0.101	0.112				
	(0.455)	(0.449)				
Geographic controls	Yes	Yes	No	No	No	No
YearXContinent dummies	Yes	Yes	No	Yes	Yes	Yes
Mean of Dep Variable	15.11	15.11	0.09	0.09	0.09	0.09
Number of cities	99	99	99	99	99	99
Number of subway cities	99	99	99	99	99	99
Number of periods	9	9	9	9	9	9
Observations	844	844	891	828	828	828

Dependent variable: Change in log region population in a 5 year period. Region exclude city itself 500km radius.

Sample is large subway cities (over 1 million in 1970). City-level clustered standard errors in parentheses.

Stars denote significance levels: * 0.10, ** 0.05, *** 0.01. All regressions $t >= 1970$.

subways quickly and that the effect is on the size of the city, not its growth rate.[19] Column 7 of table 8 investigates a somewhat different issue. Specifically, it includes leading subway measures. That these measures are significant suggests that population causes future subways rather than the converse.[20]

Finally, in table 9 we investigate whether subways have effects on the regions outside of the cities in which they are located. To begin, for each city-year define regional population for a city as total population in all other cities in our sample within 500km. In the first two columns of table 9 we report OLS regressions of a city's regional population on subways in the city, and find no effect that can be distinguished from zero. In columns 3-5 we perform first difference regressions. In column 3 our dependent variable is the change in the log of a city's regional population and the explanatory variable is the change in the log of the count of operational stations. In column four we add a control for country level GDP and continent specific year effects. We see that subways have a small positive effect on regional population growth. In column 5 we also control for the city's own population growth. The effect of subways on regional population remains positive. While these results should obviously be regarded with caution, they do suggest that cities should not build subways in order to compete with other regional cities for population.

8. Subways and urban form

The resolution of the lights data is about 1km square. This is small enough to provide some information about the way that cities are laid out, and inspection of figure 1 shows that the lights data reflect broad patterns of urban density. Given this, we use the lights data to investigate the relationship between urban centralization and subway extent. In particular, for each city-year, we calculate the intensity of light in a 5km (or 25km) disk centered on the city, and in the donut that extends from 5km to 25km from the center. We then investigate the way that light intensity in each region responds to subway extent.

Table 10 reports our results. In columns 1 and 2 our dependent variable is the ratio of total light within 5km of the center to total light within 25km of the center. As this ratio increases, a greater share of the city's light is near the center. In column 1 we conduct a pooled ols regression on our sample of 99 cities with subways in 2010 and a population above 1m in 1970. We find that centralization decreases with the count of subway stations. Column 2 conducts an analogous first differences regression and arrives at a similar conclusion. In columns 3 and 4 we conduct pooled OLS regressions. In column 3, our dependent variable is light intensity in a 5km disk and in column

[19] In an infinitely long panel, a distributed lag specification nests changes on levels and levels on levels specifications. To see this, note that if subways affected the growth rate of population we would expect all lags of a change in station capacity to have the same effect. Our finding is the opposite, only the contemporaneous subway measure matters. This indicates that change in the level of subways causes a change in the level of population but not a change in growth.

[20] If subways at t depend on population at $t-k$, then subways at $t+k$ depend on population at t. Hence, subways at $t+k$ are endogenous in a regression predicting population at t and we expect high levels of statistical significance (Ch. 10, Wooldridge (2001)).

4 it is light intensity in the surrounding donut. The effect of subways on light in both regions is positive, but is larger for the less central region. Columns 5 and 6 conduct the corresponding first difference regressions. Unsurprisingly, coefficients are smaller than in the OLS regressions, but lead to a qualitatively similar conclusion; subways increase light more rapidly in the less central region.

These results contradict the conventional wisdom that subways lead to a concentration of activity in the downtown core, and at first glance, this may seem surprising. In fact, almost any theoretical model of the spatial organization of a city will predict that the city spreads out as transportation costs fall. This is exactly what our data show. Our results are also consistent with Baum-Snow (2007), who finds that radial highways cause US cities to decentralize, and with Ahlfeldt and Wendland (2011) who find that commuter rail contributes to the decentralization of Berlin.

9. Conclusion

On the basis of figure 4, it is natural to conjecture that some sort of subway system is essential to the growth of cities beyond 5m in the rich world and that a subway system is also important to the growth of cities in the developing world.

A back of the envelope calculation also suggests that subways could have dramatic effects on the population of a city. Ten car trains can carry about 35,000 people per hour,[21] or almost 90,000 over the course of a 2.5 hour morning commute. This means that a single subway line could allow an extra 90,000 people to get to work in a central city. With a 50% labor force participation rate this leads to a population increase of 180,000. In our sample, population for a mean subway city in 2010 is about 4.7m, so this is almost a 4% population increase. Since an average subway system has 5.4 lines, adding a single line is a 19% increase. Dividing, this suggests that a theoretical upper bound for the subway elasticity of population of about 0.2. If we consider only the technical capabilities of subways, the notion that they could have an important effect on urban growth is defensible.

Our cross-sectional estimate of the effect a subway line is about twice as large as the calculation above suggests (table 4 column 6). That is, the cross-sectional estimates of the effects of subways on population are large relative to what we might reasonably guess on the basis of the physical capabilities of subways. Purely on a priori grounds, this raises the suspicion that big cities cause subways and not the converse. This suspicion finds support in our other estimates. Our first difference estimates suggest that doubling the extent of a subway network causes a tiny increase in population. While these estimates vary somewhat with technique, all are dramatically smaller than cross sectional estimates, and those we prefer are close to 1%.

To investigate the possibility that subway expansions systematically occur in years with low growth, we also conduct second difference and instrumental variables estimates. These estimates

[21] Transit Capacity and Quality of Service Manual (1999)(ch. 1, part 1, p1-22), Transit Cooperative Research Program, 2101 Constitution Ave. N.W., Washington, D.C. 20418

Table 10: Decentralization - Radiance calibrated lights

	Mean light ratio		Log mean light			
	(1)	(2)	(3)	(4)	(5)	(6)
	$lights_{5km}/lights_{25km}$	$\Delta(lights_{5km}/lights_{25km})$	$log(lights_{[0,5]km})$	$log(lights_{[5,25]km})$	$\Delta log(lights_{[0,5]km})$	$\Delta log(lights_{[5,25]km})$
$\ln(s_t)$	-0.38***		0.12***	0.24***		
	(0.091)		(0.031)	(0.040)		
$\Delta \ln(s_t)$		-0.26**			-0.020	0.040*
		(0.089)			(0.026)	(0.023)
$\ln(\text{GDP}pc_t)$	0.32		0.20**	0.11		
	(0.38)		(0.084)	(0.10)		
$\Delta \ln(\text{GDP}pc_t)$		0.41			0.77***	0.69***
		(0.29)			(0.17)	(0.12)
YearXContinent dummies	Yes	Yes	Yes	Yes	Yes	Yes
Geographic controls	Yes	No	Yes	Yes	No	No
Mean of Dep Variable	4.198	-0.293	6.139	4.654	-0.053	0.032
Mean of subways regressor	3.43	0.28	3.43	3.43	0.28	0.28
SD subways regressor	1.31	0.68	1.31	1.31	0.68	0.68
R-squared	0.34	0.30	0.41	0.50	0.52	0.50
Number of cities	99	99	99	99	99	99
Number of subway cities	99	99	99	99	99	99
Number of periods	4	3	4	4	3	3
Observations	396	297	396	396	297	297

Lights refer to mean of radiance calibrated lights in rings or circles of different diameters.

City-level robust standard errors in parentheses. Stars denote significance levels: * 0.10, ** 0.05, *** 0.01.

Geographic controls are continent dummy, capital city dummy, log km to ocean, log km to land border, and log km to major navigable river.

also yield tiny elasticities that are not statistically distinguishable from zero. Thus, the weight of evidence suggests that big cities build subways, but that subways have at most a tiny effect on urban population growth, a conclusion consistent with patterns visible in the raw data. We suspect that the similarity between instrumental variables and first differences estimates reflects the fact that, in our world sample, there is sufficient cross-country heterogeneity in the political economy of subway construction that this process is approximately random in our sample after we control for city specific effects.

While a more exhaustive analysis of the implications of subways for urban form is a subject for further research, our analysis begins this investigation. We find evidence that subways allow the central cores of large cities to spread out. This decentralization accords with the predictions of canonical theoretical models of cities: when transportation costs fall, economic activity spreads out. It is also consistent with previous analyses of the effects of radial highways on urban form. These studies also conclude that cities spread out in response to reductions in transportation costs.

It is natural to ask why the realized effects of subways on urban population diverge so dramatically from the technical frontier. Our results reflect the effects of subways that are actually built rather than their theoretical capabilities. Thus, a natural conjecture is that subways do not have much effect on city population because the subways that are actually built are not used at their full capacity. Baum-Snow and Kahn (2005) find evidence consistent with this hypothesis for the US. If true, this suggests that our results reflect a systematic failure to build useful subway systems rather than an intrinsic failure of subways to be useful. For example, our data indicate that subways are, overwhelmingly, a central city phenomenon so only people living within a few kilometers of the center can reasonably expect to walk to a station. On the other hand, much urban growth occurs on the edges of cities, e.g., Burchfield, Overman, Puga, and Turner (2006), thus subways may simply not service the areas where substantial population growth is more likely to take place. Alternatively, we know from Gordon and Willson (1984) that population density and income are good predictors of ridership. So it may be that subways are located in places where people want to live, but not where they want to ride subways.

A second conjecture also suggests itself. Consistent with within city evidence, e.g., Billings (2011) and Gibbons and Machin (2005), we find that subways reorganize activity in cities, even though they do not increase it. This suggests that an average subway is heavily used, as in our theoretical example, but rather than allowing more people to move to the city, this extra transportation capacity is used primarily to allow more travel by incumbent residents of the city. If true, this would be broadly consistent with results in Duranton and Turner (2011) on the effects of highways on travel behavior in US cities. That is, that most of the new travel caused by new highways is increased travel by incumbents. Although it is beyond the scope of the present investigation, developing a better understanding of why the observed effects of subways are so much smaller than we might predict on the basis of their physical characteristics seems like an obvious area for further research.

Many of our first difference estimates, while small, are not zero. This leads to the question of whether the effects of subways are big enough to justify a construction subsidy. To develop some intuition around this question, we suppose that a 10 percent increase in the extent of a subway system causes about a 0.1 percent increase in population. This is slightly larger than the largest of our preferred first difference estimates. It is well known that productivity increases with city size, and it is probably uncontroversial to say that city productivity increases by less than 5 percent when city size doubles. On the basis of these constants, an upper bound on the effect of a 10 percent increase in the extent of a city's subway network on aggregate city economic activity would be $0.05 \times [0.01 \times 0.1] \times 100 = 0.005\%$. Using our data on GDP and cost estimates from Baum-Snow and Kahn (2005) we can compare the value of this flow of income with the capital cost of construction. Using parameter values favorable to subway construction this calculation suggests that for an average city in our sample the value of economic activity created by subway expansion is equal to about twenty percent of the cost of construction, although the ratio of increased land rent to cost is dramatically smaller.[22] These estimates are smaller still if subways have no effect on population levels at all.

Our finding that subways have little or no effect on population growth does not seem consistent with the claims for their transformative effects sometimes made by proponents of subway construction. If subway systems have the abilities their advocates ascribe to them, then we should expect them to make cities more attractive to immigrants and hence to create population growth. Our data, therefore, broadly contradict these claims, and with them much of the justification for construction and operating subsidies. With this said, our finding does not mean that subway construction is bad public policy. Rather it suggests that the evaluation of subway projects ought to rest on the demand for mobility, farebox revenue, and not on the ability of subways to promote city growth.

[22] Calculations available on request.

Appendix

Table 11: Subways first stage: First difference – lagged subway instruments

	(1) $\Delta \ln(s_t)$	(2) $\Delta \ln(s_t)$	(3) $\Delta \ln(s_t)$	(4) $\Delta \ln(s_t)$
$\Delta \ln(s_{t-3})$	-0.085***			
	(0.018)			
$\Delta \ln(s_{t-4})$		-0.088***		
		(0.019)		
$\ln(s_{t-3})$			-0.114***	
			(0.009)	
$\ln(s_{t-4})$				-0.100***
				(0.008)
$\Delta \ln(\text{GDP}pc_t)$	0.137	0.129	-0.004	0.025
	(0.202)	(0.201)	(0.186)	(0.185)
YearXContinent dummies	Yes	Yes	Yes	Yes
Mean of Dep Variable	0.30	0.30	0.30	0.30
R-squared	0.10	0.10	0.16	0.14
Number of cities	99	99	99	99
Number of subway cities	99	99	99	99
Number of periods	9	9	9	9
Excluded instruments F-stat	21.56	21.80	145.52	153.49
Observations	828	828	828	828

Dependent variable: Change in log subway stations in a 5 year period.
Stars denote significance levels: * 0.10, ** 0.05, *** 0.01.
Sample is large subway cities (over 1 million in 1970). All regressions $t >= 1970$.
City-level clustered standard errors in parentheses.

Table 12: Population growth: heterogenous effects by size, OLS and IV

	OLS				Instrumented			
	(1) All	(2) Small	(3) Med	(4) Large	(5) All	(6) Small	(7) Med	(8) Large
$\Delta \ln(s_t)$	0.0558**	-0.0216	0.0464	0.0547*	0.0227	-0.708	0.102	0.147
	(0.0235)	(0.0910)	(0.0441)	(0.0268)	(0.145)	(0.600)	(0.185)	(0.0909)
$\Delta \ln(\text{GDP}pc_t)$	0.0978**	-0.116**	0.0887***	0.117**	0.107**	0.0000604	0.0894***	0.0656
	(0.0295)	(0.0346)	(0.0238)	(0.0538)	(0.0449)	(0.118)	(0.0219)	(0.0603)
YearXContinent dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep Variable	0.040	0.023	0.034	0.051	0.040	0.023	0.034	0.051
Mean of subways regressor	0.09	0.06	0.08	0.11	0.09	0.06	0.08	0.11
SD subways regressor	0.16	0.09	0.12	0.19	0.16	0.09	0.12	0.19
R-squared	0.39	0.67	0.39	0.51	0.38	0.02	0.38	0.45
Number of cities	75	21	26	28	75	21	26	28
Number of subway cities	75	21	26	28	75	21	26	28
Number of periods	9	9	9	9	9	9	9	9
F-stat excluded instrument					12.88	1.79	1.67	17.65
Observations	347	61	121	165	347	61	121	165

Dependent variable: Change in log population of metropolitan area in a 5 year period.

Sample is large subway cities (over 1 million in 1970) with subways systems at least 20 years old.

References

Ahlfeldt, Gabriel M. and Nicolai Wendland. 2011. Fifty years of urban accessibility: The impact of the urban railway network on the land gradient in berlin 1890-1936. *Regional Science and Urban Economics* 41: 77–88.

Arellano, Manuel and Stephen Bond. 1991. Some tests of specification for panel data: Monte carlo evidence and an application to employment. *Review of Economic Studies* 58(2): 277–297.

Barnes, Gary. 2005. The importance of trip destination in determining transit share. *Journal of Public Transportation* 8(2): 1–15.

Baum-Snow, Nathaniel. 2007. Did highways cause suburbanization? *The Quarterly Journal of Economics* 122(2): 775–805.

Baum-Snow, Nathaniel, Loren Brandt, J. Vernon Henderson, Matthew A. Turner, and Qinghua Zhang. 2012. Roads, railroads and decentralization of chinese cities. Processed, University of Toronto.

Baum-Snow, Nathaniel and Matthew E. Kahn. 2005. Effects of urban rail transit expansions: Evidence from sixteen cities, 1970-2000. *Brokings-Wharton Papers on Urban Affaires: 2005* 1(4): 147–197.

Billings, Stephen B. 2011. Estimating the value of a new transit option. *Regional Science and Urban Economics* 41(6): 525–536.

Burchfield, Marcy, Henry G. Overman, Diego Puga, and Matthew A. Turner. 2006. Causes of sprawl: A portrait from space. *Quarterly Journal of Economics* 121(2): 587–633.

Duranton, Gilles and Matthew A. Turner. 2011. The fundamental law of road congestion: Evidence from US cities. *American Economic Review* 101(6): 2616–2652.

Duranton, Gilles and Matthew A. Turner. 2012. Urban growth and transportation. *Review of Economic Studies* 79(4): 1407–1440.

Garcia-López, Miquel-Ángel. 2012. Urban spatial structure, suburbanization and transportation in Barcelona. *Journal of Urban Economics* 72: 176–190.

Garcia-López, Miquel-Ángel, Adelheid Holl, and Elisabet Viladecans-Marsal. 2013. Suburbanization and highways: When the Romans, the Bourbons and the first cars still shape Spanish cities. Universitat Autonoma de Barcelona & IEB.

Gibbons, Stephen and Stephen Machin. 2005. Valuing rail access using transport innovations. *Journal of Urban Economics* 57(1): 148–1698.

Glaeser, Edward L., Matthew E. Khan, and Jordan Rappaport. 2008. Why do the poor live in cities? the role of public transportation. *Journal of Urban Economics* 63(1): 1–24.

Gomez-Ibanez, Jose A. 1996. Big-city transit, ridership, deficits, and politics. *Journal of the American Planning Association* 62(1): 30–50.

Gordon, Peter and Richard Willson. 1984. The determinants of light-rail transit demand - an international cross-sectional comparaison. *Transportation Research Part A: General* 18(2): 135–140.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil. 2012. Measuring economic growth from outer space. *American Economic Review* 102(2): 994–1028.

Olley, G. Steven and Ariel Pakes. 1991. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64(6): 1263–1297.

Redding, Stephen J. and Matthew A. Turner. 2015. Transportation costs and the spatial organization of economic activity. In Gilles Duranton, William Strange, and J. Vernon Henderson (eds.) *Handbook of Urban and Regional Economics Volume 5*. New York: Elsevier, TBD.

Small, Kenneth A. and Erik T. Verhoef. 2007. *The Economics of Urban Transportation*. New York (NY): Routledge.

Storeygard, Adam. 2012. Farther on down the road: transport costs, trade and urban growth in Sub-Saharan Africa. Working paper, Tufts University.

Wooldridge, Jeffrey M. 2001. *Econometric Analysis of Cross Section and Panel Data*. First edition. Cambridge MA: MIT press.