STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's thesis in economics
Academic year 2018–2019

Corridors to Opportunity: A Case Study of Johannesburg, South Africa

Merve Demirel (41218) and Roshni Nandwani (41223)

Abstract

Transport equity is a powerful tool in the alleviation of transport poverty and has been gaining traction in policy discussion. Transport equity refers to the distribution of the costs and benefits of transport and the degree to which that distribution is fair or appropriate. With that in mind, this paper explores the link between public transport access and labor market outcomes in the context of South Africa. Due to the history of racial segregation, a large portion of the black population live in urban peripheries, far away from economic opportunities. In this context, we use a synthetic control method, pioneered by Abadie and Gardeazabal (2003), to examine the effect of a Bus Rapid Transit (BRT) intervention on overall and female unemployment. We find that there is a drop in unemployment that is correlated with treatment intensity, and while the average effects are not statistically significant at conventional levels, there are significant decreases in unemployment in selected years.

Keywords: Urban economics, gender, transport intervention, race, synthetic controls

JEL: C33, J61, R1, R11, R23

Supervisor: Anders Olofsgård
Date submitted: 13 May 2019
Date examined:
Discussants: Sejal Gandhi and Moritz von Schwartzenberg
Examiner: Mark Sanctuary

Acknowledgements

We would like to thank our supervisor, Anders Olofsgård, for his valuable guidance, advice, encouragement and time.

We are grateful to all scholars who have provided us with resources and feedback throughout the thesis, particularly Pamela Campa, Christian Hamann, Brightnes Risimati, Julia de Kadt, Daniel de Kadt, Lauren Graham, Ariane De Lannoy, Thea De Wet and Scott Cunningham.

We would also like to thank Erik Merkus, Harini Annadanam, Lishu Zhang and Alina Ojha for helpful discussions and input. Finally, we would like to thank our families, friends and Sebastian and Jörn for their continuous support (and excellent proof-reading).

Contents

1 Introduction **5**

2 Literature Review **7**

3 Background **10**

3.1 Transport Equity 10

3.2 Transport and Employment 11

3.2.1 Public Transport 15

3.3 South Africa 16

3.3.1 City of Johannesburg 18

3.4 Bus Rapid Transit 20

4 Data **24**

5 Methodology **25**

5.1 Synthetic Control Method 25

5.2 Our Case 26

5.2.1 Identification Challenges 26

5.2.2 Identification 27

5.2.3 The Model 30

5.2.4 Kernel Balancing 34

5.3 Implementation 34

6 Results **41**

6.1 Single Treatment Unit 41

6.2 Multiple Treatment Units 47

6.3 Kernel Balancing 53

7 Discussion **55**

7.1 Limitations 55

7.1.1 Internal Validity 55

7.1.2 Theoretical Framework 56

7.1.3 Further Research 58

8 Conclusion **59**

9 Appendix **67**

List of Figures

1 City of Johannesburg within the Gauteng province . . . 19
2 Rea Vaya phase 1a route and stops (Soweto region highlighted) . . . 22
3 Treatment intensity vs effect size . . . 37
4 Treatment effects (raw) . . . 42
5 Post-treatment p-values (raw) . . . 43
6 Treatment effects . . . 44
7 Post-treatment p-values . . . 44
8 Treatment effects, female unemployment . . . 46
9 Post-treatment p-values, female unemployment . . . 46
10 Treatment and counterfactual groups for unemployment . . . 49
11 Treatment effects for all units . . . 49
12 P-values for each year, by group . . . 50
13 Treatment and counterfactual groups for female unemployment . . . 51
14 Treatment effects for all units, for female unemployment . . . 52
15 P-values for female unemployment each year, by group . . . 52
16 Covariate balance . . . 53
17 Treated and counterfactual averages . . . 54
18 ITT for treated wards . . . 54
19 Unemployment rate - all wards . . . 68
20 Female unemployment rate - all wards . . . 68
21 Female headed households - all wards . . . 69
22 Fraction of Indian/Asian/white households - all wards . . . 69
23 City of Johannesburg electoral wards (2011) . . . 72
24 Electoral wards in Soweto (2011) . . . 73
25 Rea Vaya phase 1a route and stops with 0.7 km ward centroids (Soweto) . . . 73
26 Main places in COJ . . . 74
27 COJ with 3 km exclusion buffer of treatment . . . 75
28 Taxi network in COJ . . . 76

List of Tables

1 Predictor balance . . . 45

2 Donor weights 45
3 Predictor balance for female unemployment 47
4 Donor weights for female unemployment 47
5 ITT and significance for pooled estimates 48
6 ITT and significance for female pooled estimates 51
7 Average ITT effect for treatment wards 55
8 Predictor balance (raw) 67
9 Donor weights (raw) 67
10 Pre-treatment RMSPEs by specification (unemployment raw) 67
11 Treatment intensities by ward, kilometer distance from nearest BRT stop 67
12 Treatment effects by year 70
13 Treatment and placebo effects by year, female unemployment 70
14 Average treatment effects and significance, unemployment 71
15 Average treatment effects and significance, female unemployment 71

1 Introduction

This paper explores the role of public transport in providing individuals with access to employment opportunities, by exploring the government intervention to improve transport conditions in the Soweto neighborhood of Johannesburg, South Africa.

To grasp the context of our research, it is important to understand the concept of transport equity, which refers to the degree of fairness in the distribution of the costs and benefits of transportation (Litman, 2014). In contrast, individuals who experience a transport disadvantage face a variety of barriers that lower their degree of access to transportation (Lucas, 2012).

In the case of South Africa, transport poverty largely stems from its unique history of racial segregation; as a result of the apartheid regime, there is a high degree of physical segregation between where the disadvantaged, unemployed and mostly black populations live (townships) and where the jobs are located (city centers) (Selod and Zenou, 2001; Onatu, 2011; Allen, 2013; Kerr, 2017). Coupled with poor transport access, this has created a barrier to a perfectly competitive labor market, and the combination of individual and institutional barriers affects labor market outcomes for township residents. One of these townships is Soweto, the destination of the first forced relocation of black labor during the apartheid period. Soweto is today characterized by exceptionally high unemployment rate, low income and skill levels (Statistics South Africa, 2011), and large distances from the city center; and houses a population that is largely limited in their job search behavior and work choices. These limitations can be better understood from a rational choice perspective, where households seek to maximize utility in deciding where to live and work, or whether to seek work in another labor market (Clark and Whiteman, 1983; Monte, Redding, and Rossi-Hansberg, 2018). The cost of commuting, both in absolute terms and opportunity cost, appears in the utility function of the individual, where utility is decreasing in commuting costs. This implies that an individual is more likely to seek work or accept work if the associated commuting costs decrease (Clark and Whiteman, 1983; Monte et al., 2018). Hence, transport poverty inhibits the residents of Soweto from seeking or accepting employment in the Johannesburg city center despite the fact that business and employment opportunities largely lie in and are increasingly moving towards it.

In such a context, where inequality is spatially (and racially) distributed, public transport provision becomes crucial in shaping access to employment opportunities; particularly since marginalized households tend to rely on informal transit, spending a large share of their budget on transport, or compromising on reliability and safety for affordability (Litman, 2014), as is overwhelmingly the case in South Africa. As such, public transportation could either serve as a tool to alleviate or deepen spatial inequality (Bocarejo, Portilla, and Meléndez, 2016; Jang, An, Yi, and Lee, 2017) and public transport interventions have been gaining traction in public policy and development literature (Bhorat, Kanbur, and Mayet, 2012; Banerjee

and Sequeira, 2019).

In line with this, the South African government implemented a Bus Rapid Transit (BRT) system, named 'Rea Vaya', in order to meet the City of Johannesburg's increasingly dire and growing transport needs. The BRT is an express bus service that functions like an above-ground metro, running on lanes dedicated to BRT buses. It is characterized as being affordable, efficient and convenient, and hence offering a high degree of urban mobility (Onatu, 2011; Allen, 2013). The Rea Vaya is also highly accessible, with the trunk, feeder and complementary routes ensuring that most Soweto residents are within walking distance from the nearest bus stop (compiled using data from Risimati and Gumbo, 2018 and Statistics South Africa, 2019b).

Consequently, as the BRT reduces both the absolute and opportunity cost of travel by providing a cheaper and shorter, congestion-free commute, we expect the implementation to push Soweto residents towards job search and job acceptance in the city as rational choice theory dictates, resulting in more of the previously unemployed job-seekers now working in Johannesburg.

At the same time, this framework leaves the potential of gender-differentiated outcomes; particularly since there is empirical evidence that commuting costs have a much heavier weight in women's work choices, leading them to compromise on work opportunities in favor of shorter, safer and cheaper commuting (Fanning Madden, 1981; Fagnani, 1987; Bank, 2011). This effect is more pronounced for women in marginalized neighbourhoods. In this context, the BRT system is believed to be particularly beneficial for women as stations are manned and monitored while the areas around the minibus taxi pick-up zones are often unsafe (Onatu, 2011; Allen, 2013).

As such, from a policy perspective, it is essential to understand if BRT systems like the Rea Vaya are viable in bridging the gaps in transport provision for the socially disadvantaged. In particular, we seek to understand if the Rea Vaya's goals of addressing historical inequalities in the post-apartheid city (Allen, 2013) are in fact being achieved, and if there are heterogenous effects on subgroups of the Soweto population.

Our paper examines the effect of BRT access on the unemployment rate in Soweto, as well as on the female unemployment rate, using data from the South African census and Quality of Life Surveys on an administrative ward level. In addition, we examine further heterogenous effects based on the distance from BRT stops. We use a Synthetic Control Method, an extension to difference-in-difference estimation; as opposed to taking a single control unit or an average of multiple control units, this method constructs a data-driven synthetic control unit using an optimally-weighted average of all control units. The use of synthetic controls has been increasingly popular in case studies, and to quote Athey and Imbens, 'is arguably the most important innovation in the policy evaluation literature in the last 15 years' (Athey and Imbens, 2017). This method is particularly fitting in our case, especially given the identification restrictions.

We begin with an overview of the literature in Section 2, followed by a detailed conceptual and

contextual background in Section 3. Section 4 describes the data sources, while in Section 5 we detail the identification strategy. Section 6 and 7 present the results, discussion and implications for future research, and we end with concluding remarks in Section 8.

2 Literature Review

Transport literature is one that is truly interesting to study; the relationship between social disadvantage, public transport and labor market access is founded on the intersection of various strands of literature including urban economics, trade, development, economic geography and labor economics. Stemming from such a diverse collection of research, yet bringing them together in a novel manner, our research question builds on a small but growing portion of economic literature that focuses on evaluating the long-term outcomes of specific public transport interventions. We proceed to briefly discuss how our research question is related to these broader research efforts, however we significantly limit our discussion to areas most relevant to our research in an attempt to be concise.

One strand of the broader literature consists of studies related to investments in transport infrastructure, such as the quality and provision of roads, the introduction of extended railway networks and the connection of rural areas to cities, or the effect of these on trade opportunities in terms of providing access to markets and consequently on output and productivity. This literature makes use of the gravity equation from trade theory, and as such is closely related to trade literature. Donaldson and Hornbeck (2016) examine the historical impact of railroads on the agricultural sector of the American economy in 1890 and find an increase in county market access, with a consequential increase in agricultural land value as the railroad network expands. Donaldson (2018) also studies the welfare effects of the railroads constructed in colonial India, and finds that the railroad network reduced trade costs and increased real income levels. Laird and Venables (2017) create an appraisal framework to measure the economic impact of transport investment in a city, arguing that it brings comparative advantage and economies of scale by opening channels for trade, creates employment opportunities through job market access, and ensures that skilled labor is geographically concentrated. Melo, Graham, and Brage-Ardao (2013) conduct an empirical meta-analysis on the productivity effects of transport infrastructure and find varying effects across industries, continents and modes of transport.

On the other hand, urban economics literature focuses on residential choices, housing markets and their relation to labor markets, and studies how these are affected by spatial and geographical disadvantages. Brueckner and Zenou (2003) examine the effect of racial discrimination in suburban housing markets in the US that forces black workers to live in cities, far away from suburban employment, and show that both employment and wages are lower for blacks living in the city center than in the suburbs. On the other hand,

Glaeser, Kahn, and Rappaport (2008) find that central cities disproportionately attract the poor in the U.S., mainly due to the greater access to public transport they provide. This finding is important because it highlights the shortcomings of the traditional housing market theory in explaining the sorting of the poor into inner cities; with transportation mode choice having three to four times the explanatory power of income elasticity of land, the implication is that the public transport dependent poor sort into denser inner cities precisely because of their dependency on transport.

More specifically, the effect of commuting time and commuting cost have been studied extensively in urban economics literature; and forms the primary motivation for our research. In this area, researchers have focused on identifying the effects of commuting, both on the demand and supply side of labor markets. Most of these studies have also explored different mechanisms of discrimination in the labor market; Carlsson, Reshid, and Rooth (2017) study the individual labor market effects of living in deprived neighborhoods in Sweden, through discrimination in hiring practices. While they find no evidence of general neighborhood signaling effects as measured by interview callback rates, applicants with commuting times longer than 90 minutes have lower callback rates regardless of their neighborhood of residence. Zax and Kain (1996) exploit the exogeneity of a firm's relocation from Detroit's central city to a white suburb, and find that while white employees move to get closer to the firm's new location, the increase in commuting distance causes many black workers to quit their jobs. Gutiérrez-i-Puigarnau and van Ommeren (2010) analyze the relationship between labor supply and commuting distance using socioeconomic panel data for Germany between 1997 and 2007 with a similar empirical strategy. In contrast to Zax and Kain (1996), they do not find a negative effect of commuting distance on total labor supply, suggesting that significant disparities exist in different country contexts.

Zooming in on public transport papers, a majority of the literature is focused on measuring primary effects, specifically accessibility effects. Accessibility can be measured using a variety of metrics such as utility-based measures (opportunity cost of transport expense and time), activity-based measures (number of jobs or any other opportunities within reach for a given commute time) or infrastructure-based measures (number of people that can be transported at any given point) (Geurs, van Eck, et al., 2003). This work has demonstrated the role for impact evaluation of transport policy; such as the Jang et al. (2017) study of the distribution of spatial inequality in Seoul and how it is affected by a new subway line. The authors discuss the improvement to the public transport system on the intensive and extensive margin and conclude that while an improvement to the frequency of transport mode deepens spatial inequality, an extension of the transport network would reduce it. Similarly, Bocarejo et al. (2016) study the effect of the BRT implementation in Bogota on social fragmentation, measuring the degree of social interaction across income groups. A difference-in-difference approach between the three different income groups shows that the BRT and the feeder routes, have indeed reduced social fragmentation.

In one of the few studies conducted in a developing country setting, Majid, Malik, and Vyborny (2018) look at the effect of the BRT system in Lahore on accessibility and other outcomes. Using a matching approach and distance to the nearest stop as an instrument for accessibility (measured by commuting time and cost to a central point in the city), they also measure spillover effects (the commuting times for other vehicle types), the number of vehicles owned, the proportion of the population working outside the home, years of education and migratory effects. The authors find that access to the BRT reduced both the time and cost of commuting, caused workers to switch from private to public modes of transport, increased public transport use and attracted more highly educated people to the neighborhood. Picarelli, Jaupart, Chen, and Moreno-Monroy (2017) provide a detailed overview of the spatial inequities in Accra, Ghana by different commuting modes, and identify a non-causal link where better access to employment opportunities is correlated with a higher likelihood of employment and a lower probability of engaging in the informal sector.

Finally, there is a small pool of literature looking at the secondary outcomes of public transport provision. Zimran (2018) studies the impact of transportation on health in rural U.S. and finds that greater transportation linkage, measured by the level of market access, had a negative impact on health outcomes by increasing population densities. Although sparse, there exists some empirical literature on the relationship between public transport and labor market opportunities. Holzer, Quigley, and Raphael (2003) look at the effect of the opening of San Francisco BART rail stations on the individual employment opportunities of minority workers, and find that accessibility 'matters slightly for Latino workers but not for African American workers'. Tsivanidis (2018) finds an increase in accessibility from the BRT implementation in Bogota, Columbia, and that this increase in access to jobs is associated with an increase in wages. In addition, he finds an increase in commuting distances which was greater for low-skilled workers, suggesting that 'the system made employment in more distant locations more attractive' (Tsivanidis, 2018). Rotger and Nielsen (2015) investigate the effect of increased public transit access in Copenhagen on individual earnings, via a new metro route connecting suburbs to the city center, and find a positive effect. They attribute this partly to a 'change in commuting patterns as the improved access to public transport facilitates a shift from employment within the township to better paid jobs in the city centre' (Rotger and Nielsen, 2015). Our paper fits right in with this strand of literature, and is one of the very few focusing on the evaluation of a public transport intervention in a developing country.

3 Background

3.1 Transport Equity

Transport infrastructure and transport services are an important component of development policy. On a macroeconomic level, transport is an important determinant of economic interaction and market access. It allows for the movement of goods and people, for similar economic activity to cluster and for cities to develop comparative advantage. This in turn increases the productivity of economic activity and transportation affects the spatial reallocation of the benefits associated with the increase (Melo et al., 2013; Laird and Venables, 2017;). On a microeconomic level, transport is crucial in linking individuals to economic and social life, ensuring their participation in the broader society and access to goods and services such as education, healthcare and leisure facilities. For the rest of this study, we focus specifically on the microeconomic role of transport and explore the concept of 'transport equity'.

Transport equity has been studied since as early as 1973, when Wachs and Kumagai identified individual mobility as an important component in addressing inequality (Wachs and Kumagai, 1973). It refers to the distribution of the costs and benefits of transport and the degree to which that distribution is fair or appropriate (Litman, 2014). Consequently, 'transport disadvantage' is defined as the situation where an individual has difficulty accessing transport, driven by factors such as high cost, poor transport availability, safety concerns and a lack of information on public transport (Lucas, 2012). The specific ways in which an individual can be excluded from accessing transport systems are namely geographical exclusion (where an individual is unable to access transport services due to their housing location), economic exclusion (where the high cost of transport limits access to travel options), physical exclusion (where physical barriers such as vehicle design or the lack of disabled facilities prevents travel), time-based exclusion (where an individual's own time constraints such as household and childcare responsibilities reduce the time available for travel), space exclusion (where access to space on transport is reduced, for example in differentiated train compartments) and fear-based exclusion (where personal safety concerns prevent the use of transport services) (Church, Frost, and Sullivan, 2000).

Transport disadvantage amplifies social disadvantage; communities characterized by factors such as low income, low employment levels, low skills or poor housing become further entrenched in 'transport poverty'. Transport poverty is a phenomenon characterized by an inaccessibility to goods and services, opportunities and social networks as well as an exclusion from decision-making and planning processes driven by the interaction of the two forces (Church et al., 2000; Levitas et al., 2007; Lucas, 2012). A social environment built around the assumption of high mobility further traps the individual in a cycle of deepening transport and social disadvantage (Kenyon, 2006). In this sense, mobility and redistribution of transport wealth can be linked to several dimensions of poverty, including the inability to accumulate social and

human capital, as well as insecurity and vulnerability in contexts where geographical proximity is prominent (**booth˙poverty˙nodate**). As such, transport poverty research is less concerned with the absolute provision of transport within a community but rather the implications of transport deprivation with regards to an individual's access to 'life-enhancing opportunities such as employment, education, healthcare and their supporting social networks' (Lucas, 2012).

There is vast evidence in the literature of the link between social disadvantage and transport disadvantage, where significant inequalities exist in the access to transport and travel behavior of marginalized groups. Johnson, Currie, and Stanley (2011) find that low income households in Australia have difficulty accessing employment, education and leisure due to transport constraints, while Altman and Hinkson (2007) and Lucas (2011) find a similar effect for aboriginal populations and black South Africans respectively. It is also widely documented that the transport burden falls disproportionately on poor households and women (Fanning Madden, 1981; Fagnani, 1987; Bank, 2011). Supporting these findings, in development literature we see evidence of positive effects of mobility improvements on development outcomes. Dercon and Hoddinott (2005) find that improving the quality and presence of roads and transportation in rural Ethiopian villages significantly and substantially increases consumption and decreases poverty. The authors also find that communities with access to better roads have persistently higher growth rates than others (Dercon and Hoddinott, 2005). In addition, there is evidence indicating that children go to school more often if they live closer to it; Vuri (2010) shows that reducing the commuting distances to schools increases school attendance and reduces household chore activities. This study focuses on labor market outcomes.

3.2 Transport and Employment

While there is a broader literature that focuses on the social and economic effects of transport disadvantage and transport poverty, we draw attention to sections of the literature that incorporate transport into a rational choice perspective and examine employment outcomes. This literature assumes that households make trade-offs between affordable housing and transport, and a lack of provision of either results in the household having to forego safety, accessibility or amenities (Litman, 2014).

The mechanisms that link underemployment and transport poverty are based on the combination of individual restrictions with a wider overall context of poor job access (van Ham, Mulder, and Hooimeijer, 2001). Therefore, to understand how transport poverty affects underemployment, it is important to make a distinction between these mechanisms and circumstances that exacerbate already adverse outcomes, such as gendered and racial disadvantages. We first briefly discuss and describe key mechanisms, followed by a closer look on how gender and racial dimensions might paint a more complicated picture.

Poor geographical access to jobs can cause adverse labor market outcomes through both supply

and demand side mechanisms. On the supply side, long and expensive commutes are a deterrent to the effective hours and days worked, as well as the take-up of work that is too far away, given the wages offered. Secondly, spatial disconnect from jobs result in higher job search costs that might cause workers to restrict their search to an inefficiently small spatial limit, not search as intensively or search with decreased efficiency. On the demand side, firms might be unwilling to hire workers that live further away, refusing to pay the extra cost associated with getting that worker to work, or citing lower productivity from longer commutes. As a result, spatially constrained or spatially inflexible individuals become constrained in their employment choices (Gobillon, Selod, and Zenou, 2007). We examine two key theoretical models that further elaborate on supply-side mechanisms.

Theory 1: Clark and Whiteman (1983) examine the job-searching behavior of spatially constrained low-income workers, arguing that the microeconomics of the rich and poor are the same, where they both make rational decisions based on their preferences. However, the environment in which these decisions are made results in differences in wages, job quality and job location. In particular, they assert that 'in a world of inequality, rational behavior can reinforce disequilibrium and the perpetuation of spatial-economic inequality' and conclude that urban poverty is often not a result of poor individual motivation, and where the poor choose to live is not due to any kind of 'sub-optimal social attachment to a particular community', but rather that external barriers are an important determinant of individual behavior (Clark and Whiteman, 1983).

The Clark and Whiteman (1983) paper builds a model in which workers decide on whether or not to undertake a job search in a labor market that is geographically distinct from where they live, focusing on market conditions that inhibit labor market adjustment in the short run and create a spatially segmented labor market between two locations in the same metropolitan area. The model assumes that travel expenses are borne by the worker (as is often the case in low-wage jobs) and includes foregone income in the search cost parameter. In their utility-maximizing model, present utility is a function of initial wealth, wages and past job search costs. In the event of an unsuccessful job search, utility is lower than the present level as search costs are incurred with all other variables (such as wages) remaining unchanged. On the other hand, the expected utility of a successful job search (receiving and accepting a new job) is a function of the wage distribution of all firms in the searched market, the likelihood of the offered wage being above the worker's reservation wage, and the search costs incurred. This implies that if search costs are high enough compared to the expected wages, conducting a search in a wide spatial area might result in greater incurred costs than the benefits one expects to receive, and hence a lower utility than that of the present.

The authors argue that this is indeed a possibility for low-wage, low-skill workers searching for employment located far away, exacerbated by the fact that demand for low-skilled labor is often much smaller than supply and hence the worker spends a longer time (and incurs more costs) in the job search. As a result, the poor tend to be less ambitious in their job seeking behavior. At the same time, according to the

model, the worker constructs an expected utility of the search using the probabilities of success and failure and only conducts the search if this expected utility is greater than his or her present utility; this is highly influenced by search costs that enter both functions. As such, public transport improvements can serve to reduce job search costs in multiple ways. While fare reductions directly reduce these costs, time saved in travel reduces search costs by reducing the foregone income. As mentioned above, this can be significant given the wage levels and the unfair amount of time that low-skilled workers have to spend in searching. As such, by increasing the expected utility of both a successful and unsuccessful search, improvements in transport lead to a larger differential between expected and present utility. This implies that workers who were previously below the threshold might be pushed over the boundary, causing them to undertake a job search. The discussion is easily extendable to the case of unemployed workers.

Theory 2: Monte et al. (2018) develop a general equilibrium model of migration and commuting decisions, incorporating spatial linkages between locations in both the goods and labor markets to determine the local employment elasticity. One aspect of the model focuses on the factors that cause an individual to live and work at specific localities. Similar to Clark and Whiteman (1983), the model computes commuting costs as both transport costs and the foregone income from commuting.

In the model, workers choose their residential and work locations to maximize utility, which is a function of wages and the commuting costs between the home and workplace, amongst other parameters. The authors determine that all else equal, workers are more likely to work in location i and live in location n if the wages in i increase or if commuting costs decease. The authors go on to determine that the conditional probability that a worker living in location n works in location i is a function of the wage-commuting cost ratio of working in i as a fraction of the sum of wage-commuting cost ratios of all possible work locations, i.e. the worker is most likely to take up work where the wage-commuting cost ratio is the largest, all else equal.

The model is found to be consistent with empirical findings (McFadden, 1974; Grogger and Hanson, 2011; Kennan and Walker, 2011) implying that the wage-commuting cost ratio is an important determinant of whether a worker is likely to take up work that involves an expensive commute. In the case of low-wage, low-skill workers, this ratio is expected to be small as commuting costs can constitute a large fraction of wages. This emphasizes the role of public transport in increasing the radius of employment opportunities that spatially-constrained individuals are willing to take up, hence bridging the gap of transport disadvantage.

Moving on to the additional factors that are widely identified to exacerbate transport poverty, in the following sections we examine and provide evidence of these phenomena in the context of race and gender disadvantage.

Race The Spatial Mismatch Hypothesis (SMH), put forward by Kain (1968), argues that the spatial mismatch between the areas where the appropriate skill force lived and the areas where appropriate jobs

were located was the main driving force of the Watts riots in Los Angeles, as it led to adverse labor market outcomes for black citizens. Among the low-income, low-skilled populations living in the city center, those who were able to relocate could, in theory, also decentralize towards the suburbs.

While this is in line with conventional economic theory that dictates that utility-maximizing households choose their residential location with respect to access to jobs, this population was predominantly black and was facing racial discrimination in the housing market. Hence, while the jobs were moving towards the suburbs due to lower rents, black households were not able to move to where the opportunities were or afford to commute the long distances, resulting in adverse labor market outcomes specifically for their community (Kain, 1968).

While the original SMH focuses on the specific case of disconnectivity from job opportunities of low-skilled minority workers in city centers, it does an excellent job illustrating how such mismatch can result in involuntary unemployment for any spatially inflexible individual. Another example of how discrimination affects job search behavior is a study by Abel (2017) in South Africa that finds suggestive evidence that immigrants, facing racial discrimination in their home suburbs, extend their job search to other areas. This adjusting of job search behavior means that the worker incurs an additional cost of discrimination in the form of transport cost in searching and working in places further away.

Gender Women face different social and time constraints compared to men, and as such, the two theories previously described are much more pronounced for spatially constrained women. As a result, commuting circumstances (safety, cost, time and quality of transport mode) have a much heavier weight in women's work and residential choices. This usually means that women disproportionately compromise on the quality of opportunities in favor of shorter, safer and cheaper commuting.

Specifically, one widely documented empirical fact is that women have shorter commuting distances and times than men (Gordon, Kumar, and Richardson, 1989; Blumen, 1994; McLafferty, 1997; Turner and Niemeier, 1997). The reasons why women do not commute far are likely to be a combination of the factors mentioned above, and there is a growing body of work focused on laying out these mechanisms. Gendered household responsibilities play a significant role in limiting women's work options; shouldering a disproportionate amount of household burden, women are time-constrained and tend to accept jobs closer to home (Fanning Madden, 1981). A survey of mothers in Marseille (Fagnani, 1987) found that they 'preferred to take any kind of job near their residence rather than commute further away to get a better-paid and higher-status job'. A study in India showed that women are willing to compromise on quality of education for safety in the public transport used, and this willingness to pay for safety translates to a 20 percent decline in their post-college wages (Borker, 2017). An obvious consequence of this trade-off is that spatially inflexible women, 'who often reside in more marginal neighborhoods where most available jobs are informal and low in

productivity', face greater effects on the quality of work (Bank, 2011). As such, interventions that aim to increase spatially inflexible/immobile women's access to economic opportunities should consider the unique set of challenges women face that amplify the low access to jobs.

3.2.1 Public Transport

Urban communities that have the least geographical access to social and economic opportunities and simultaneously face higher mobility constraints are particularly dependent on the provision and quality of public transportation. In the absence of formal public service provision, the choice households face is between owning a car and informal paratransit, which usually results in households spending a larger share of their budget on transport, or compromising the reliability and safety of transportation modes for affordability (Litman, 2014). Therefore, public transport provision is especially important in shaping access to employment when inequality is spatially distributed, and interventions focused on public transport subsidies have been increasingly receiving attention (Bhorat et al., 2012; Banerjee and Sequeira, 2019).

We focus our discussion on public transport. Consequently, the next logical question to ask is what constitutes good public transport. Gomide, Leite, and Rebelo (2005) evaluate the adequacy of public transport services using 4 broad categories, namely affordability (the monthly spending required with respect to a user's income, opportunity costs of the service and the provision of benefits or subsidies), availability (route coverage, waiting times, journey times, the availability of services on weekends and at night, as well as reliability and regularity), accessibility (the walking distance to stations, availability of information, access roads and the necessary environment for the disabled or elderly users) and acceptability (the level of safety, quality of service and capacity). Out of these, transport fares, the distance to stations, waiting times, journey times, the level of delays and irregularities in service provision, crowd levels and safety have been identified as especially relevant for low-income populations (Gomide et al., 2005). Hence it is with these factors in mind that we develop our research question and empirical model.

In this setting, we argue that there is scope for policy intervention, as it is concluded throughout the literature that where you live matters, and that the spatially inflexible are usually dependent on affordable transport to mitigate the effects of residential constraints. As such, transport policy has the potential to mitigate not only the adverse effects of transport disadvantage, but also to correct for the inequalities in other dimensions like race and gender. Economic theory suggests that interventions focusing on commuting costs can level the playing field. In such a context, public transportation could either serve as a tool to alleviate or deepen spatial inequality, and the effects of such interventions should be evaluated rigorously (Bocarejo et al., 2016; Jang et al., 2017).

3.3 South Africa

Employment The political history of South Africa has had a significant impact on the shaping of its socioeconomic landscape, leading to a high degree of inequality in income, employment, human capital and access to opportunities between blacks and whites (Selod and Zenou, 2001) that still persist today.

South Africa has one of the highest unemployment rates in the world, at 25 percent in 2014 as compared to that of the OECD countries, of 7 percent (OECD, 2019a). One of the key drivers is the strikingly high youth unemployment, at 51 percent in 2014 versus the OECD average of 15 percent (OECD, 2019b). Unsurprisingly, this is largely segmented along racial lines. In 1994, the unemployment rate of blacks was 26 percent, reaching a peak of 37 percent in 2002 and 27 percent in 2007. In contrast, the unemployment rate of whites remained at 4 percent in 1994 and 2007, peaking at 6 percent in 2002 (Mlatsheni and Leibbrandt, 2014). While 56 percent of South Africans lived below the poverty line in 2015, the poor were disproportionately black, with 93 percent of poor persons being black despite the racial group making up only 75 percent of the population. Blacks also earn 4.3 times lower wages on average than whites and 1.9 times lower wages than 'coloureds' (the second lowest income group in the country) (BusinessTech, 2015).

These disparities and social inequities extend to education levels and consequently human capital; according to the 2001 census, only 22 percent of blacks have completed high school (known as 'matric') as compared to 71 percent of whites (Statistics South Africa, 2004). While these education outcomes are post-apartheid, they have been found to be rooted in apartheid laws, where education for blacks was segregated to low-quality schooling, evidenced by a large pupil per teacher ratio in 1991 of 42 for blacks (versus 19 for whites) and a disproportionately small share of government education expenditure for blacks of 47 percent (Selod and Zenou, 2001).

Segregation The apartheid system that was in place from 1948 to 1994 also highly influenced spatial planning, with institutionalized racial segregation leading to the spatial segregation of residential and commercial land areas that had huge ramifications on the access to opportunities for black residents (Onatu, 2011; Allen, 2013). As part of the restrictions on land-use, only whites were allowed to live near city centers, where most employment opportunities were located. Blacks and other non-white groups (coloureds and asians/indians) were only allowed to live in the periphery of cities, where race-specific townships were created under the Group Areas Act, separated by buffer zones to prevent cross-community interactions (Selod and Zenou, 2001; Kerr, 2017). Such a design was aimed at keeping the African labour pool sufficiently distant from economic and white areas but close enough that the labour could be used for low-paying jobs (Allen, 2013).

The racially motivated land use policies were repealed in 1994, and while some desegregation was observed in the late 1980s, these spatial segregation patterns remain visible in the present day. As of 2007, 55

percent of black adults were still living in townships, and 81 percent of township residents had reported their intention to continue living there for various reasons such as community affiliations, economic opportunities in the townships, or the inability to afford alternative housing locations (Findley and Ogbu, 2011). This phenomenon is in line with Selod and Zenou (2001), who model post-apartheid residential decisions and propose that 'after apartheid, the structure of the city remains the same; black families reside in the outskirts of the city', while 'white families locate in the vicinity of the city center' (Selod and Zenou, 2001).

As a result, South African cities today have very contrasting urban zones. In Cape Town for example, the central city area hosts a concentration of middle and high-wage jobs while the peripheral townships are inhabited by the middle to lower-income with poor access to services (Selod and Zenou, 2001). With most townships being 25 to 30 kilometres away from the city's central business district, low-income black communities are also far away from industrial and commercial work opportunities in economic centres, with a large degree of separation between the workplace and residence (Selod and Zenou, 2001; Onatu, 2011; Allen, 2013; Kerr, 2017).

Commuting and Transport Consequently, these spatial challenges result in significant commuting and transport challenges for the low-income black community, who face high monetary and time costs of commuting due to their residential locations (Allen, 2013; Kerr, 2017). According to Onatu (2011), low income earners in South Africa spend over 10 percent of their income on transport, while the average commuting distance to work for blacks (over 15 kilometers) is more than twice as long as for whites (less than 7 kilometers) and the average daily commuting time is 50 percent longer (Selod and Zenou, 2001; Statistics South Africa, 2014).

This situation is made worse by the fact that the country suffers from one of the longest commute times and worst congestion levels in the world (Kerr, 2017). A study by IBM in 2011 constructed a Commuter Pain Index on the world's most unfriendly cities for commutes and ranked Johannesburg in fifth place, after Beijing, Shenzhen and Mexico City (Gyimesi, Vincent, and Lamba, 2011), while the 2011 OECD's 'How's Life' study ranked South Africa in first place for the highest average commute time of 56 minutes per day compared to a 24-country average of 38 minutes (OECD, 2011). Between 1993 and 2003, the mean real cost of transport increased from 261 to 460 Rands (Kerr, 2017). One of the consequences of the high commuting costs is that the poor spend the bulk of their energy walking or waiting for public transport, with transport safety a huge concern as traffic accidents are frequent with largely pedestrian fatalities (Onatu, 2011). Blacks are also much more public-transport dependent than whites, with 65 percent reporting using public transport to get to work as compared to 3 percent of whites in 2003 (Kerr, 2017).

3.3.1 City of Johannesburg

Unemployment Being the country's economic capital and most densely urbanized area (Allen, 2013), we focus our attention on the Gauteng province, consisting of several municipalities including the City of Johannesburg (Figure 1). The City of Johannesburg consists of a bustling city center surrounded by suburbs and townships, and hosts large economic opportunity and attractive employment. Between 2001 and 2011, the municipality was the most preferred destination for migrants, attracting 38 percent of all migrants and 49 percent of foreign migrants to the country, a phenomenon explained by Ngobeni (2014) as attributable to the attraction of higher earning potential, as Gauteng province was one of the two provinces with a higher income share than its population share in 2001 (Ngobeni, 2014). In addition, Rogerson and Rogerson (1997) find that small businesses owned by blacks were moving towards the city center, further increasing the economic potential of the city center.

The City of Johannesburg is comprised of a township on its south-western periphery known as Soweto. Soweto was the destination of the first forced relocation of black labor pools from rural to urban areas during the apartheid period, and was formed in 1949 when 80,000 Africans were relocated there. As of the 2011 census, close to 20 percent of Soweto residents had no income, with approximately 60 percent earning below 38,000 Rand a year (Statistics South Africa, 2011).

As such, despite a growing economy, the City of Johannesburg faces extreme income disparity (Allen, 2013), and mirrors the contrasting urban landscape seen in South African cities, where black townships like Soweto are spatially segregated from economic opportunities in the city center.

Current transport situation in Soweto While the affluent rely on private cars for transportation, the intentional lack of investment in effective and efficient public transport to the townships, coupled with the large physical separation between the home and workplace for non-white communities led to the development of an informal public transit sector. This bottom-up system served the purpose of filling the service gap between the need for urban transport and the capacity of the state system, however without state regulation or support (Findley and Ogbu, 2011; Kumar, Zimmerman, and Agarwal, 2012; Allen, 2013). In particular, 'minibus taxis' developed to cover popular routes between the home and workplace, operating in an informal market where commuters could get onto a shared vehicle heading towards their destination (Allen, 2013). Initially, the minibus taxis were viewed as an entrepreneurial trend, where one could easily enter the market by loaning or purchasing a vehicle (Allen, 2013). This was especially relevant for black populations with lower access to the formal labor market (Harrison et al., 2010).

However, the combination of low investments and high returns led to exponential growth in the industry and there was soon an number of uncoordinated private operators in the market hoping to make

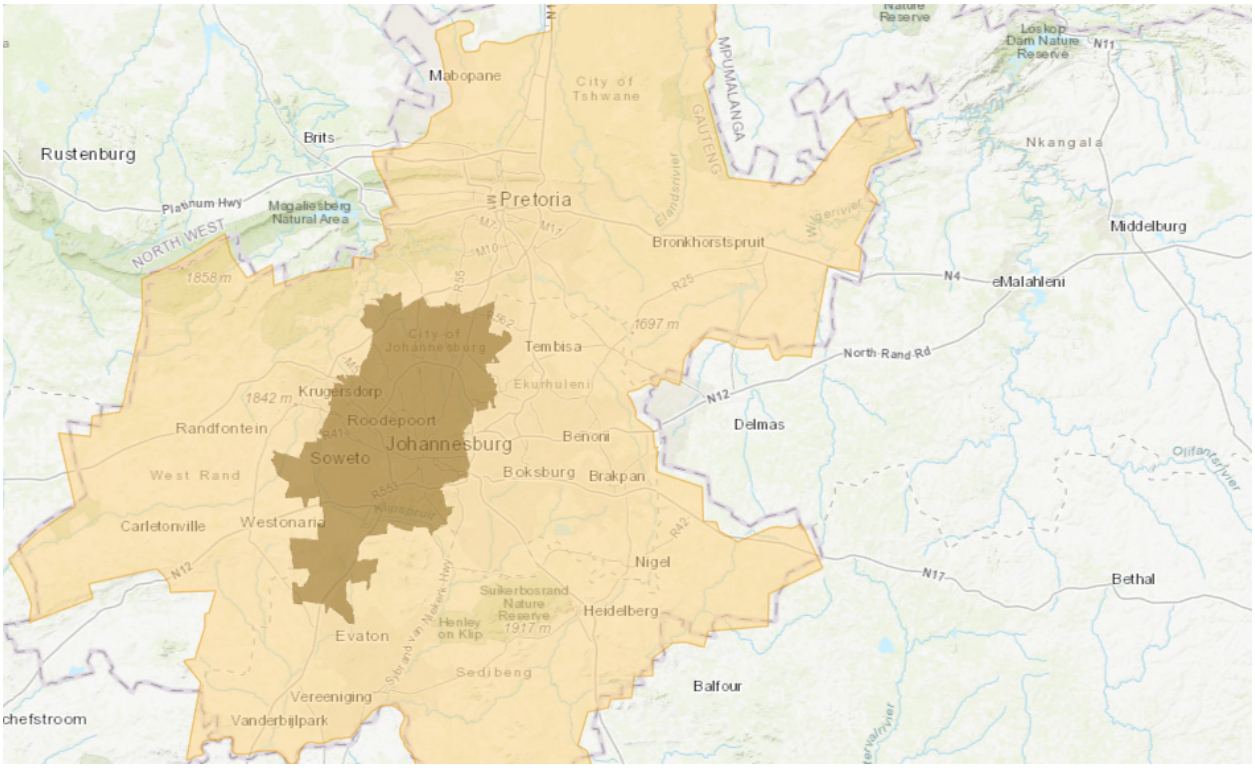


Figure 1: City of Johannesburg within the Gauteng province
Source: Authors' illustration of Municipal Demarcation Board maps

financial gain (Allen, 2013; Woolf and Joubert, 2013). By 2003, the minibus taxis were the most frequented mode of public transport, covering 72 percent of all public transport trips in Gauteng in 2003, while only 14 percent and 9 percent were covered by trains and city buses respectively (Allen, 2013). In 2013, there were 32 taxi associations controlling over a thousand different routes (in both directions), operating from over 450 different starting points (Johannesburg Strategic ITP Framework, 2013).

By the 1990s, the emergence of private players to meet the dire transport need, without regulation, led to a system that provided the best return to the owners, while marginalizing the user with long travel times, high and unstable fares and dangerously maintained vehicles driven by drivers with poor skills (Kumar et al., 2012; Allen, 2013). According to Woolf and Joubert (2013), the market remained informal and unsubsidized, with taxi operators receiving no formal training or employment protection. The lack of integrated fares among the players led to volatile, unstable pricing and as a result, households spent a disproportionately high proportion of their budget on transport, a phenomenon that was more pronounced in poorer black households that had longer commutes from the townships (Allen, 2013). According to Lucas (2011), the key issue with transport and consequently with access to services was not the absence of public transport but rather that they were 'unaffordable, unsafe, unreliable and unsuitable' for the long journeys that black commuters often undertake to go to work or other critical social institutions (Lucas, 2011). In addition, these informal transport services contributed heavily to congestion, air pollution and road accidents (Venter,

Jennings, Hidalgo, and Valderrama Pineda, 2018).

Similarly, another form of public transport, the metered taxis, are reported to provide low-quality service. Even though several formal large companies are present in the market, about 44 percent of the metered taxis are 'no-name' taxis that are not regulated or monitored. It is common for metered taxi vehicles to be operated in a state of poor repair, employ drivers with poor knowledge of their area of operation, and operate without meters, without radio or GPS control or operating licenses. Most importantly, this resulted in fluctuating fares and unsafe transportation for users (Johannesburg Strategic ITP Framework, 2013).

As such, there were huge disparities in access to transport, with residents in the townships facing two to three hour commutes on informal transit or walking while affluent families had two or even three cars (Findley and Ogbu, 2011; Woolf and Joubert, 2013; Machard, 2014). According to Machard (2014), the significant commuting costs that township residents faced in both the monetary sense and in the time spent in accessing, waiting and travelling on public transport significantly reduced their potential for undertaking other crucial, 'life-supporting' activities (Machard, 2014). Lucas (2011) identifies four factors that contribute to how the already socially disadvantaged poor in the South African townships suffer from transport poverty and consequently poor access to economic opportunities; the limited access to private transport that drives a reliance on public transit, the prevalent minibus taxis that have proven to be unreliable and inaccessible, the apartheid-based spatial planning and urban growth that exacerbated commuting costs, and the reliance of the poor on walking, which leads to a high degree of exposure to air pollution and road hazards. In this context, public transport provision is all the more crucial in addressing transport disadvantage; with a history of spatial segregation, the majority of blacks living in urban areas still live in townships on the periphery of cities (and intend to continue doing so), far away from jobs as the townships were built to be inaccessible from the city. This is exacerbated by the social disadvantage the black community faces, with predominantly low income and employment and a high level of dependence on public transport that consumes a large share of their time and income.

3.4 Bus Rapid Transit

Background BRT systems are a particular type of bus networks that are predominantly implemented in cities with 'high levels of inequality, informality, and low spatial quality' (Venter et al., 2018). The reason is that BRT systems are positioned as being 'pro-poor', aimed at alleviating poverty by improving accessibility and reducing transport costs and time (Venter et al., 2018). The World Bank has financed several BRT initiatives in recent years based on this argument (Gilbert, 2008) and since its conception, BRTs have been implemented in over 40 locations worldwide with another 80 planned to be underway (Onatu, 2011). The Rea Vaya, a BRT system implemented in Gauteng, is an express bus service that functions like an above-ground metro, with the trunk routes running on dedicated lanes that run down the center of existing inner city

roads, decreasing the commuting time lost to congestion. The buses run on a high frequency of up to three minutes during peak hours, providing access to and within the city for 19 hours a day as opposed to the demand-driven informal system that is peak-hour heavy. The buses are able to accommodate large passenger volumes with a seating capacity of up to 112 people on the trunk route that covers over 300 kilometers and has weather-proof bus stations every 500 metres. The buses also allow for kerbside boarding outside of bus lanes, and has passing lanes for express buses, multiple stopping bays and off-board ticket purchasing to accommodate high passenger traffic (Onatu, 2011; Allen, 2013). In addition to being affordable, efficient and convenient, and hence offering a high degree of urban mobility, the system is particularly beneficial to women as stations are manned and monitored whereas areas around minibus pick-up zones were often unsafe (Onatu, 2011; Allen, 2013). The Rea Vaya also addresses the congestion issue at large, reducing the number of vehicles on overburdened roads and the resulting carbon emissions (Onatu, 2011).

Rollout and Implementation There are two main drivers of the Rea Vaya implementation as cited in the literature. Firstly, there was a need to provide a formal and regulated alternative to the minibus taxi industry to address the lack of an integrated and accessible public transport system that served the needs of the less affluent (Woolf and Joubert, 2013; Machard, 2014). Secondly, the awarding of the FIFA World Cup event to South Africa meant that there was both increased funding and need to establish a mass transport system that could accommodate the surge in visitors and facilitate their movement to the World Cup events. Twelve cities, including Johannesburg, were chosen to receive support for the upgrading and integration of all public transport modes (Allen, 2013). At the same time, the decision to implement a BRT system amongst other public transport alternatives stemmed from the observed success of BRTs in other developing countries like Colombia and Brazil (Allen, 2013; Woolf and Joubert, 2013). As such, the government decided to develop a BRT system adapted to South Africa in the hopes of formalizing public transit in the city and improve transport access for all user classes (Allen, 2013; Woolf and Joubert, 2013; Machard, 2014).

The BRT system in South Africa was the first full BRT to be implemented on the African continent. It was named 'Rea Vaya' and was developed in phases, with the first phase (phase 1A) being brought into operation in stages between 31 August 2009 and February 2011[1] (Johannesburg Strategic ITP Framework, 2013). Figure 2 shows the route of phase 1A, consisting of a 25.5 kilometer dedicated roadway that connected Thokoza Park in Soweto to Ellis Park on the east side of the Johannesburg central business district, which was a popular minibus taxi route (Kumar et al., 2012; Allen, 2013; Woolf and Joubert, 2013). The route comprises of a single main 'trunk route', with two 'complementary' and five 'feeder' routes that augment the trunk by connecting commuters from smaller neighborhoods (Kumar et al., 2012; Allen, 2013). The implementation was not without tribulation, however, as it faced stiff resistance and violent clashes with the strong taxi unions, despite efforts to consult and involve them during the planning phase the awarding of

[1]Phase 1B of Rea Vaya, rolled-out in late 2013, is outside the scope of this paper. The additional stops were outside of Soweto, therefore does not affect the treatment status of our sample.

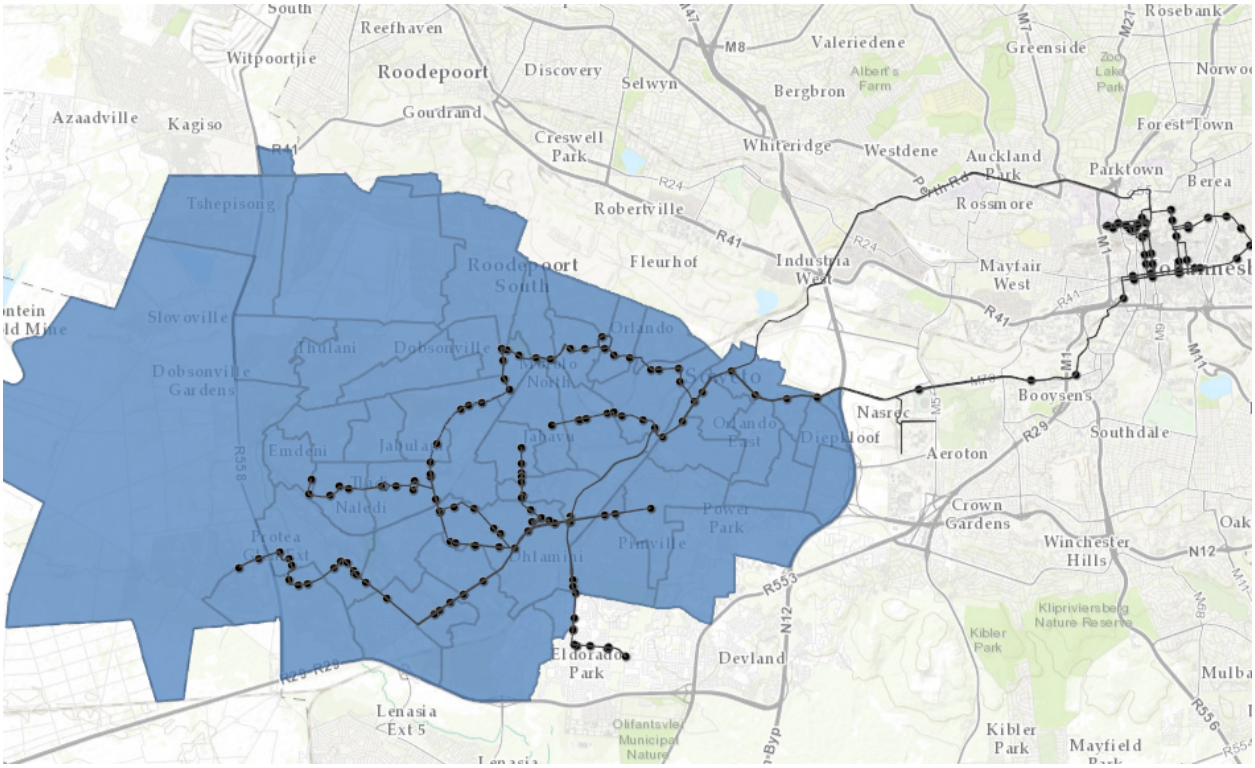


Figure 2: Rea Vaya phase 1a route and stops (Soweto region highlighted)
Source: Authors' illustration of Municipal Demarcation Board maps

minibus taxi operators with shares in the system and offering minibus taxi drivers employment as Rea Vaya drivers (Onatu, 2011; Kumar et al., 2012; Allen, 2013; Woolf and Joubert, 2013).

Motivation and Research Question A review of the literature reveals that multiple sources have concluded that the Rea Vaya was an unsuccessful initiative, with some calling it a 'mammoth failure' (Venter and Hayes, 2017). This is often attributed to a 'low take-up rate', where the government is far from being able to recover the millions invested and operating costs (Scorcia and Munoz-Raskin, 2019). However, in determining whether the Rea Vaya was a failure of the system in providing viable transport, it is important the note that (1) this opinion is reflective of an investment shortfall rather than the quality and characteristics of the system, and (2) this shortfall should be evaluated in its unique context. According to Scorcia and Munoz-Raskin (2019), the history of spatial segregation in the urban areas of Johannesburg highly influenced travel patterns, leading to long, unidirectional trips with very high passenger volume in peak hours versus very low numbers in non-peak times, as opposed to typical city travel patterns of shorter journeys with higher turnovers that are necessary for mass transit systems to be financially sustainable. As such, the Rea Vaya might not be able to cover operating costs from fare box revenue simply because of its operating conditions.

However, it is important to differentiate the welfare aspects from the business case for such an

intervention; with its unique features, the Rea Vaya system is largely in line with the previously highlighted definition of good public transport. Although information on government fare subsidies is scarce, the fares are frequently cited as being affordable to the average Soweto resident (Onatu, 2011; Allen, 2013). The 'availability' criteria, that includes journey times, waiting times, regularity and operating times, is largely fulfilled with the high frequency buses operating for a major part of the day on dedicated roadway to avoid traffic congestion. The Rea Vaya is also highly accessible, with the trunk, feeder and complementary routes ensuring that most Soweto residents are within walking distance from the nearest bus stop (compiled using data from Risimati and Gumbo, 2018 and Statistics South Africa, 2019b). With high security and manned stations, the Rea Vaya also scores highly on the fourth criteria of 'acceptability'.

This is reflected in the growing ridership numbers. According to Kumar et al. (2012) and Allen (2013), total daily ridership in the first phase of implementation was 17,000 by December 2010 and grew to 45,000 in mid-2011. In addition, about 60 percent of Rea Vaya users in Johannesburg reported a saving in travel cost of about 21 percent (Vaz and Venter, 2012). As such, from a policy perspective, it is still important to understand if BRT systems like the Rea Vaya are viable in bridging the gaps in transport provision for the socially disadvantaged. In particular, we seek to understand if the Rea Vaya's goals of addressing historical inequalities in the post-apartheid city (Allen, 2013) are in fact being achieved, and if there are heterogenous effects on sub-groups of the Soweto population. Within this context, our paper examines the effect of access to the Rea Vaya on the unemployment rate in Soweto, as well as on the female unemployment rate. In addition, we examine heterogenous effects based on the distance to the nearest Rea Vaya stop.

Why we expect a significant effect Based on contextual and theoretical evidence, we expect a significant effect. South Africa's unique history of racial segregation has been a major driver of transport poverty, and the physical segregation between where the unemployed live (townships) and where the jobs are (city centers) is a barrier to a perfectly competitive labor market. This combination of individual and institutional barriers affects labor market outcomes through the mechanisms outlined in the previous section. With the exceptionally high unemployment rate, low income and skill levels in Soweto and the large distances between the home and city center, we expect a large proportion of the population to lie below the threshold of the job searching conditions outlined in Theory 1 and the threshold of wage-commuting cost ratios in Theory 2. Hence these residents do not seek or accept employment in the Johannesburg city center despite the fact that business and employment opportunities largely lie in and are increasingly moving towards it. This is corroborated by Bhorat (2012), who argues that in this context, it is 'extremely expensive for unemployed individuals to go out and search for jobs, as the jobs exist in areas which are extremely expensive to get to' and that 'the cost of searching for work is unaffordable for these individuals' (Bhorat, 2012). De Lannoy, Graham, Patel, and Leibbrandt (2018) find that the average job search cost (made up mostly of transport

costs) in a sample of young people participating in youth employability programmes is almost double of their per capita monthly household income, and they often resort to borrowing money from friends and family to look for work (De Lannoy et al., 2018). Consequently, as the BRT reduces both absolute travel cost and foregone income by providing a cheaper and shorter, congestion-free commute, we expect the implementation to push Soweto residents above these thresholds and make both job search and job acceptance in the city a rational choice, resulting in more of the previously unemployed job seekers now working in Johannesburg.

4 Data

We use quality of life (QOL) surveys by the Gauteng City-Region Observatory (GCRO) (Gauteng City-Region Observatory, 2019a; Gauteng City-Region Observatory, 2019b; Gauteng City-Region Observatory, 2019c; Gauteng City-Region Observatory, 2019d; Gauteng City-Region Observatory, 2019e), supplemented by census data from the Statistics Agency of South Africa (StatsSA) (Statistics South Africa, 2019a) for our analysis.

The GCRO conducted five waves of QOL surveys across the Gauteng province, in the years 2009, 2011, 2013, 2015 and 2018. The survey collects data on the demographics of the enumerated population (population group, gender, age, language), housing (dwelling type, tenure, satisfaction with dwelling, perceived quality of housing and housing allocation), household services (water, sanitation, refuse, energy sources), migration, health (including disability), education and employment (including employment sector), community services and amenities, transport, leisure activities, safety and crime, financial data, perceived personal well-being and quality of life. The QOL is available on an electoral ward-level of aggregation, which is an administrative unit smaller than a province, a district municipality or local municipality, and is constructed according to the number of voters in a given region (Municipal Demarcation Board, 2016). The QOL datasets also contain weights for each observation to ensure a representative sample that mimics the ward's proportion in the population census. These weights are specific to each ward-gender-population group combination and were applied as suggested.

The census data from years 1996, 2001 and 2011 provided by StatsSA captures basic descriptive and labor market characteristics as well as details on household amenities such as running water, electricity, toilet facilities and dwellings. Even though QOL surveys have a considerably richer collection of data, our analysis is limited to indicators available across both datasets for consistency. As with the QOL surveys, census data is available on an electoral-ward level.

We also have geolocation data on roads, bus and train routes and stops (including the Rea Vaya), taxi routes, business districts and informal housing districts in Gauteng, obtained from Risimati and Gumbo

(2018), geolocation data on electoral ward maps, obtained from the GCRO and Municipal Demarcation Board (Municipal Demarcation Board, 2019; Observatory, 2019), and on main place, obtained from Statistics South Africa (2019b).

5 Methodology

Our goal is to determine the causal impact of the improved public transport access on the population of Soweto, specifically whether it results in increased levels of employment. We do this by following a synthetic control approach (SCM), as outlined by Abadie and Gardeazabal (2003), Abadie, Diamond, and Hainmueller (2010), and Abadie, Diamond, and Hainmueller (2015). We apply this method first for a single treatment unit, then for multiple treatment units using disaggregated data. We also provide simple trajectory balancing estimates, an extension of SCM suitable for mitigating data limitations developed by Hazlett and Xu (2018).

5.1 Synthetic Control Method

The idea of a synthetic control is to construct an appropriate counterfactual for the treated unit, i.e. what would have happened to the path of employment of the affected neighborhoods in the absence of the BRT line, and to assess the intervention's impact by comparing the counterfactual to the actual path observed. SCM constructs the counterfactual unit (a 'synthetic' Soweto) by using optimally weighted untreated units from the pool of untreated ('donor') units. The estimator is then the difference in the post-treatment outcomes between the treated and counterfactual unit.

The synthetic control estimator is essentially an extension of linear panel data methods, a factor based model allowing the effects of unobserved confounders to vary with time (Abadie et al., 2010). In studies where a difference-in-difference estimation would traditionally be employed, one has to assume that the trajectory of the main outcome variable between the treatment and control group is similar, so that the pre-treatment trends can be extrapolated to form the estimate. However, in reality it is often difficult to find a control unit that adequately resembles our treatment unit, excepting for the treatment status, or whether the parallel trends assumption holds. It might also be difficult to measure the similarities between the treatment and control unit and assess whether they are adequate or not. Instead, the SCM offers a data-driven method of constructing appropriate control groups in a transparent way. Since a synthetic control, or synthetically generated counterfactual, is a weighted average of the available donor units, the SCM makes explicit the relative contribution of each donor unit to the counterfactual as well as the similarities between the treated unit and the synthetic control in terms of pre-intervention outcomes and other predictors of post-intervention outcomes (Abadie et al., 2010). As such, the researcher can evaluate how well the synthetic

control's outcomes match the treated unit's outcomes before the policy change, and the similarities between both groups (McClelland, 2017). Constructing the control group as a weighted-average of many donor units also allows multiple time-varying confounders, meaning that the parallel trends assumption is not required. This makes the SCM especially well-suited for examining the effects of a population-level treatment, i.e. case studies where aggregate data is available and control units with parallel trends might not be available.

Since the original paper, there have been many variants of the synthetic control estimator including synthetic difference-in-difference, augmented SCM, generalized synthetic control estimator as well as extensions that accommodate multiple treatment units, continuous treatment status, multiple treatment periods and interactive fixed effects (Gobillon and Magnac, 2013; Doudchenko and Imbens, 2016; Powell, 2017; Xu, 2017; Ben-Michael, Feller, and Rothstein, 2018). Below we present the original model and discuss the relevant variations that we implement - multiple-treatment units and kernel balancing.

5.2 Our Case

5.2.1 Identification Challenges

In this section, we discuss why SCM is a fitting solution to our case. In an ideal scenario, we would examine the structural relationship between accessibility and labor market outcomes in two stages; first quantifying the overall accessibility changes from the public transport intervention, and then determining the causal impact of this improved accessibility on labor market outcomes.

Our main identification problems are three-fold;

- **Non-random allocation of treatment:** The neighborhoods that received the intervention were not randomly allocated. Therefore, it is not possible to identify the effect of this intervention simply by comparing it to other parts of the City of Johannesburg, since there might be systematic differences in those particular neighborhoods that were treated. In other words, our intervention is not exogenous to the dependent variable. In Section 5.2.2, we review the transport literature to examine how this issue has been dealt with by other researchers, and how best to navigate this source of endogeneity.
- **Combining multiple data sources:** We rely on two data sources, one census, the other a representative survey. Reliability of the latter will always be imperfect due to the nature of sampling. We discuss the comparability of these two sources in Section 6.
- **Reliance on aggregate data:** We only observe aggregate-level outcomes, and since we do not track a panel of individuals throughout the time period, we are limited to reduced form estimation of labor market outcomes. We cannot examine whether the intervention increased individual levels of

accessibility and improved labor market outcomes in relation to the increase in accessibility. We are not able to distinguish sorting effects, nor the mechanisms that were previously discussed in Section 3.2 that might lead to improved labor market outcomes. Since we only track wards across time, we are only able to focus on aggregate outcomes, treating this intervention as a case study. Researchers have typically studied this relationship with the aid of general equilibrium models and novel datasets, see Tsivanidis (2018) and Donaldson (2018) for two rigorous studies jointly examining housing market, commuting and labor market decisions. Even for a much less comprehensive analysis, this would require us to track a group of individuals, collecting data on their commuting distances and lengths, distance to the nearest public transport station, average expenditure on commuting or job search, as well as availability, location and characteristics of employment across time. Facing these data constraints, we limit our study to the reduced form relationship while recognizing the areas of improvement and avenues for future research. Moreover, our data does not allow us to distinguish any reorganization of the economic activity around the new BRT line from the causal effect of the intervention itself. Although this is a possible source of reverse causality, we believe the reorganization effects to be minimal since the intervention was implemented on an already existing transport corridor.

In addition, case studies are usually conducted on policy changes that are expected to be a major explanatory factor in the outcome variable. While we recognize that the explanatory power of public transport interventions might not be as high as what is conventionally studied, given the lack of literature quantifying the effects of public transport interventions, we are convinced that it is an important area of study.

It is with these limitations in mind that our identification strategy is designed, and while the nature of our research question and available data restricts our capacity to address them more rigorously, we limit their effects where possible.

5.2.2 Identification

In the broad literature identifying the causal impact of transportation, instrumental variables have been influential, namely in the use of historical route IVs, planned route IVs, and inconsequential place approach. Redding and Turner (2015) provide a review of the studies using these approaches; most of which focus on estimating the causal impact of transportation networks on the economic outcomes of a region, rather than a particular transport line. These instruments could also be used as placebo tests, as Donaldson (2018) does with a planned route IV in estimating the effects of railroads in India. While these methods are popular, they are not suitable as the main identification strategy for our research question. Randomized experiments and matching strategies have been less widely applied; we are not aware of any studies using matching strategies in our particular context and the nature of randomization makes it difficult to conduct such studies in not just transport but in overall infrastructure investments.

Within public transport provision, the literature has heavily focused on modelling the first-stage accessibility changes or correlations with potential outcomes (Picarelli et al., 2017; Majid et al., 2018). In the extension of the literature that establishes a causal relationship to employment outcomes, a common method to overcome non-random assignment of a transport line is to use difference-in-difference estimation, by having a distance-to-station cutoff line to divide the observations around the public transport stations into treatment and control groups. This relies on the assumption that the two groups will be very similar in characteristics except for their treatment status; take-up rate will be much higher for those living closer, and for the ones living further away it will not be a major change in accessibility therefore low-to-no take-up. By mimicking a random assignment of treatment, the difference in the outcome variable of interest between the two groups could be credibly attributed to the intervention. This methodology is preferable if there is micro-data present, such as in Rotger and Nielsen (2015), that quantify the effects of a new metro station in Copenhagen.

However, a simple difference-in-difference strategy is not well-suited for our case, since not having a micro-panel makes it difficult to credibly assign treatment and control status to observations. If we follow the approach of assigning a cutoff radius within Soweto, there is a high risk of spillover effects as we do not observe where the BRT users live. As such, we cannot be sure of which areas will take the treatment, and which ones will not. We are further discouraged from treating any wards within Soweto as control units upon discovering, from our datasets, that it is not uncommon for the Soweto population to walk over 30 minutes (approximately 3 kilometers) to their first mode of public transport, which covers almost the entire Soweto region. If we instead choose the cut-off line as Soweto itself, we risk having treatment and control groups that follow different paths, particularly due to the unique urbanization structure of Johannesburg where townships are high in population density, but relatively far apart from each other.

Having exhausted the approaches in the literature, we turn to alternative methods of identification. SCM is particularly well-suited for our case, where only aggregate data is available and many time-varying confounders affecting unemployment are likely to exist. However, our limited pre-treatment data becomes a particular limitation when implementing the SCM. We discuss the fit of an SCM approach to our research question and data, and compare SCM with traditional difference-in-difference and matching below.

SCM vs Difference-in-difference A traditional difference-in-difference attempts to control for the differential trends across treatment and control groups by controlling for unit specific trends in the econometric model. This involves the extrapolation of pre-shock trends to post-shock periods, which is a strong assumption, especially over long periods of time. While it is possible to test for pre-existing trends through event study analysis, it is not generally clear how one should proceed if this test fails (Cavallo, Galiani, Noy, and Pantano, 2012).

In light of this, the typical motivation for the use of SCM is that in reality, the 'parallel trends' assumption required by additive fixed effect models often does not hold. SCM permits the 'estimation of specifications with interactive fixed effects which generate non-parallel trends, permitting causal inference even when additive fixed effect models are inappropriate' (Powell, 2017). The ability of the SCM method to relax the parallel trends assumption is important for our identification strategy; we only observe three pre-treatment time periods, with large gaps in between them (1996, 2001 and 2009), and are unable to credibly claim parallel trends with any control group we choose.

Even though the SCM makes it possible for us to estimate a counterfactual, it does not solve our data limitations entirely, and as such we exercise caution as we implement the methodology and examine the results of the SCM matching closely. We discuss the choice of matching strategy and predictor variables in Section 5.3.

SCM vs Matching Matching is another method, usually employed with micro-data with many treated units, that estimates the missing potential outcome for each treatment unit using a control group unit that is 'closest' to the treatment unit for some X confounder. This gives us estimates of all the counterfactuals from which we could simply take the average over the differences. However, matching for more than one covariate creates a curse of dimensionality, especially when the sample size is small (as is common in case studies where we have one treated unit and few non-treated donor units) (Cunningham, 2018). Moreover, the SCM allows non-equal weighting of the donor units in constructing the counterfactual, in contrast to conventional regression methods, generating a closer match for the counterfactual (Dube and Zipperer, 2015). The SCM is therefore a better fit for assessing the impact of a treatment implemented in one or very few units.

The advantages of the SCM vis-à-vis matching is less straightforward when there are many treatment and donor units. A unique synthetic control requires that the values of the matching variables for a treated unit fall in the convex hull of the corresponding values for the donor pool. In cases where there are many donor units, there might be multiple possible convex combinations of untreated units that perfectly reproduce the values of the matching variables for the treated observation, i.e. the optimally-weighted synthetic control might not be unique. This curse of dimensionality can typically be resolved by ad-hoc methods, like increasing the number of covariates or by restricting the donor pool to units that are similar to the treated units such as was done by Acemoglu, Johnson, Kermani, Kwak, and Mitton (2016). In such a case, the researcher 'faces a trade-off between minimizing the covariate discrepancy between each treated unit and its synthetic control as a whole (pure synthetic control case) and minimizing the covariate discrepancy between each treated unit and each unit that contributes to its synthetic control (pure matching case)' (L'Hour and Abadie, 2018). The literature is, at the moment, inconclusive on the best practices on generalization of this issue; such as on how to decide on the maximum number of units that contribute to the synthetic controls or on how to

proceed when the synthetic controls are constructed with multiple characteristics. L'Hour and Abadie (2018) suggest a penalized estimator as a solution in their working paper, however they have indicated that the work is incomplete at the moment. Taking these limitations into account, we conclude that the SCM offers advantages that are well-suited for our research question.

5.2.3 The Model

The SCM model by Abadie et al. (2010) is as follows:

Let Y_{it}^N be the outcome that would be observed in unit i at time t in the absence of the intervention, for units $i = 1, ..., J+1$, and time periods $t = 1, ..., T$. Let T_0 be the number of pre-intervention periods, with $1 \leq T_0 \leq T$. Let Y_{it}^I be the outcome that would be observed for unit i at time t if unit i receives the intervention in periods $T_0, T_1 ...$ to T. We assume that the intervention has no effect on the outcome in the pre-treatment period, hence for all $t\epsilon\{1, ..., T_0\}$ and all $i\epsilon\{1, ..., N\}$, we have that $Y_i^1 = Y_{it}^N$.

Let $\alpha_{it} = Y_{it}^I - Y_{it}^N$ be the effect of the treatment on unit i at time t, and let D_{it} be a treatment indicator that takes the value 1 if unit i receives the intervention at time t, and value 0 otherwise. The observed outcome for unit i at time t is

$$Y_{it} = Y_{it}^N + \alpha_t D_{it} \tag{1}$$

Since only the first unit, unit 1, receives the intervention and only after period T_0 where $(1 \leq T_0 < T)$, we have that $D_{it} = 1$ if $i = 1$ and $t > T_0$, 0 otherwise. Note that unit 1 is the treated unit while all other units are part of the 'donor pool', or units that may potentially compose part of the synthetic control. We aim to estimate $(\alpha_{T_0+1}^i, ..., \alpha_T^1)$. For $t > T_0$, we have that

$$\alpha_{1t} = Y_{1t}^I Y_{1t}^N = Y_t^1 Y_{1t}^N \tag{2}$$

Since Y_{1t}^I is observed, in order to estimate α_t^1 we simply need to estimate Y_{1t}^N. Note that implicit in the model is the assumption that the donor units do not receive treatment throughout the observed period, and the treatment in the treated units does not affect the outcomes in the donor units.

Suppose that Y^N is given by a factor model

$$Y_{it}^N = \delta_t + \theta_t Z_i + \lambda_t \mu_{it} + \epsilon_{it} \tag{3}$$

where

- δ_t is an unknown common factor with constant factor loadings across units,

- Z_i is a $(r \times 1)$ vector of observed covariates (not affected by the intervention),
- θ_t is a $(1 \times r)$ vector of unknown parameters,
- λ_t is a $(1 \times F)$ vector of unobserved common factors,
- μ_i is an $(F \times 1)$ vector of unknown factor loadings and
- ϵ_{it} error term represents unobserved transitory shocks at the unit level with zero mean.

We note here that equation (1) becomes the usual difference-in-difference (fixed-effects) model if we impose that μ_t is constant in equation (1) for all t. In other words, the difference-in-difference model allows for the presence of unobserved confounders but restricts their effects to be constant over time, hence by taking time difference, they can be eliminated. In contrast to this, the factor model in the SCM allows for the effects of confounding unobserved characteristics to vary with time. In this model, taking the time differences does not eliminate the unobserved confounders, μ_j.

The SCM aims to find a weighted combination of units W, in the donor pool such that the following condition holds:

$$\sum_{j}^{J+1} = (w_j^* Y_j) = Y_1 \tag{4}$$

where w_j represents the weight on unit j. If this condition holds, then $\sum_{j=2}^{J+1}(w_j^* Y_{jt})$ provides an estimate of Y_{1t} when $D_{1t} = 0$. This is done by solving for the optimal donor and predictor weights (denoted by matrices W^* and V^*) that minimize the distance ($||X_1 X_0 W||_V$) between the vector of pre-intervention characteristics of the treated units and donor units (X_1 and X_0, respectively), where

$$(||X_1 X_0 W||_V) = \sqrt{(X_1 X_0 W)' V (X_1 X_0 W)} \tag{5}$$

and V is some k x k symmetric, positive semidefinite matrix.

The feasibility of solving for a weight matrix W^* such that this condition holds depends on whether the trajectory of treatment unit can be approximated with some linear combination of donor units. Namely, equation (4) will hold exactly only if $(Y_{11}, ..., Y_{2T_0}, Z_2')$ belongs to the convex hull of the corresponding sequences for the donor units. One possibility is to restrict the pool of donor units to those with X_1 values similar to the X_0 values of the treated unit. The method is valid for any choice of V^*, leaving the researcher with room to influence the matching. However, note that the choice of predictor variables included in it will affect the quality of the match between treatment and synthetic control units. Ideally, the predictor variables should be chosen such that the relationship between the outcome and predictor variable is linear, in order to avoid possible interpolation bias.

Given that we are able to come up with weights such that this condition holds, the estimator for the treatment effect becomes:

$$\hat{\alpha_{1t}} = Y_{1t} - \sum_{j=2}^{J+1}(w_j^* Y_{jt}) \tag{6}$$

Extension to Multiple Treated Units Synthetic controls have become increasingly popular with comparative case studies; Cavallo et al. (2012) first implemented the method to identify the effect of natural disasters on mid-term growth, followed by applications to minimum wage increases such as Dube and Zipperer (2015), and the use of disaggragated data such as in Gobillon and Magnac (2013) and Acemoglu et al. (2016). The motivation to apply SCM for multiple treated units could either simply be to examine the treatment effect with a large number of treated and control units, to compare and contrast events that are otherwise not directly comparable and reach a more generalized empirical conclusion, or to prevent interpolation biases by using disaggregated data (L'Hour and Abadie, 2018). As an extension to the initial model, we use disaggragated data in order to examine our results in finer detail.

The SCM lends itself to multiple treatment units easily, as the model described above could be generalized to incorporate multiple treatment units in the dataset. Pooling the treatment effects could be done in a variety of ways, one being as simple as taking the mean of treatment effects across treated units, as introduced by Cavallo et al. (2012) and as we will later on use. Recall that our period-specific estimates of the treatment effect on the treated units are denoted by $(\hat{\alpha}_{1,T_{0+1}}, ..., \hat{\alpha}_{1,T})$ for periods $1, 2, ..., T - T_0$. Then the estimated average effect for the G treatments is given by

$$\bar{\alpha} = (\bar{\alpha}_{g,T_{0+1}}, ..., \bar{\alpha}_{g,T}) = \frac{1}{G}\sum_{g=1}^{G}(\alpha_{g,T_{0+1}}, ..., \alpha_{g,T}) \tag{7}$$

Inference Inference in SCM is less formal compared to other methods. Abadie et al. (2010) suggest a placebo test similar to permutation inference, where the distribution of a test statistic is computed under random permutations of the sample units' assignments to the intervention and nonintervention groups. The SCM is applied to every potential control in the sample, which allows us to calculate the exact distribution of estimated effects from the placebo interventions. Under the assumption of randomized intervention, the exact p-value can be calculated; in cases where the intervention is not random, the p-values for the treatment effect can be interpreted as 'the possibility of the affect occurring purely by chance' (Abadie et al., 2010).

To put it more formally, this inference approach involves (1) iteratively applying the SCM to each unit in the donor pool and obtaining a distribution of placebo effects, (2) calculating post-treatment root mean squared prediction error (RMSPE) ratio (or post-to-pre treatment RMSPE ratio) where the RMSPE

for each period is computed as follows:

$$RMSPE = \sqrt{\frac{1}{T - T_0} \sum_{t=T_0+t}^{T} \left(Y_{1t} - \sum_{j=2}^{J=1} w_j Y_{jt} \right)^2} \tag{8}$$

and the test statistic to be used is the treatment unit's position in the distribution of this ratio in descending order:

$$p = \frac{RANK - 1}{TOTAL - 1} \tag{9}$$

This exact p-value can then be compared to conventional levels of statistical significance (Cavallo et al., 2012).

A straightforward extension of permutation inference, rank-sum inference, has been discussed by Dube and Zipperer (2015) in the context of pooling multiple treatment effects. The authors show that the mean of the percentile ranks of the treatment effects can be used to judge the statistical significance for the pooled estimate. The test statistic for each unit is defined as follows:

$$p_g = \hat{F}_g(\alpha_{g1}) \tag{10}$$

where $\hat{F}_g$ is the CDF of treatment effects, and

$$\bar{p} = \frac{\sum_{g=1}^{G} p_g}{G} \tag{11}$$

is the test statistic for pooled estimates. Under the assumption that the percentile rank is uniformly distributed, statistical significance of the individual treatment effects can be assessed in a similar manner as before. Distribution of the pooled estimate can be calculated using the Irwing-Hall distribution as the sum of G independent uniform random variables, and converges to a normal distribution by the Central Limit Theorem; allowing us to once more infer statistical significance with conventional 5 percent levels.

As mentioned before, the inference approaches developed in the literature are not free of issues; namely, since the donor pool stays constant throughout the inference process, synthetic controls are likely to include a similar combination of these donor units. Therefore, the strong assumption that the treatment and placebo effects, and consequently RMSPE ratios are independent and identically distributed need not hold. See Dube and Zipperer (2015), Firpo and Possebom (2016) and L'Hour and Abadie (2018) for a more

detailed discussion.

5.2.4 Kernel Balancing

We extend our analysis to mitigate some of our data limitations by employing a less data intensive 'kernel balancing' method developed by Hazlett and Xu (2018). More broadly, the trajectory balancing method is a general re-weighting approach to causal inference with time-series cross-sectional data. The specific variant of trajectory balancing we implement, kernel balancing, is different to our methodology only in the sense that it seeks for a set of weights that matches not only on the mean trajectory of the outcome variable, but on higher order of statistical characteristics such as variance or volatility. The intuition here is that with large T_0, there are more constraints imposed on the optimization, therefore more factors that have the potential to influence the post-treatment outcome is captured in the synthetic control. Taking from Hazlett and Xu (2018), 'in the absence of long pre-treatment periods, a control group that varies wildly around a flat line could be well mean-balanced to a treated group that has all 'flat' trajectories by giving equal weights to each control unit. If features like variance or volatility later come to have a large directional impact on Y_0, then this imbalance can generate bias'. By introducing such features on the matching criteria, kernel balancing mimics the ability of synthetic controls to tolerate time-varying confounders, but takes advantage of few pre-treatment constraints to impose stricter constraints.

To illustrate simply, let $K(Y_i, Y_j)$ be a function that measures similarity of Y_i and Y_j. With N observations, we can construct an (NxN) kernel matrix K such that $K_{ij} = k(Y_i, Y_j)$. The method then proposes that the pre-treatment outcome variable to be matched is replaced with this vector K, or ϕY_i^{pre}, containing data assessing how similar unit i is to unit 1, unit 2, and so on. Since an exact balance on all N features is likely to be infeasible, the matching is done on the first P principal components of K, which could mean that the matching will be done only on mean and variance.

5.3 Implementation

We conduct our analysis in three steps:

1. Single treatment unit (Soweto)
2. Multiple treatment units (41 wards within Soweto)
3. Trajectory balancing approach

Single Treatment Unit This level of analysis takes Soweto as a treatment unit, allowing us to estimate treatment effects on the entire township level. As one of the goals of the Rea Vaya was to provide Soweto residents with better access to the city center, taking the whole township as a treated unit would provide an understanding of how close it has come to achieving these goals, by allowing us to determine the overall effect of the intervention on Soweto residents.

This definition of the treatment unit is in line with our proposed definition of the treatment group as 'households that are at most 3 kilometers away from the closest Rea Vaya stop'. This is based on our previously mentioned observation that it is not uncommon for the Soweto population to walk over 30 minutes to their first mode of public transport, that translates to approximately 3 kilometers. With this definition, the entire Soweto township lies within the treatment group. Hence we construct a 'synthetic Soweto' to model counterfactual outcomes from a weighted average of a number of donor units.

Another reason for this level of treatment is that it allows us to overcome limitations with the available data. In order to implement the synthetic control method, we need to be able to track the changes in the outcome and predictor variables on a consistent geographical level. The smallest aggregation level that we are able to obtain data for is the electoral ward. One inconsistency in these datasets is that the boundaries of Gauteng wards change in 2011, implying that a ward-level comparison is not feasible across years. Although StatsSA provided all census data with the 2011 ward boundaries, the 2009 QOL survey is only available on the pre-2011 boundaries. In order to overcome this problem, we have the possibility to use a smaller-area disaggregation provided in the 2013 QOL wave and onwards, called the Enumeration Area or Small Area Layer (Lehohla, 2003), to reconstruct pre-2011 ward-level data for all QOL datasets. However, we do not have access to this level of aggregation for the census datasets and are hence unable to reconstruct it for census data. Since 1996, 2001 and 2009 are the only pre-trend data we have, we cannot drop the census datasets and therefore this approach is not feasible to use with the synthetic control method. Instead, we aggregate the ward level data to a higher level, constructing a dataset on a 'main place' level that has consistent boundaries across the years. Main place is level five of seven in the geographical structure constructed by StatsSA, defined as 'an easily recognizable, small-area geographical entity at local level, and corresponds, for example, with the name of the city, town, tribal area or administrative area' (Lehohla, 2003). Within this structure, each element in each level can be geo-coded to identify it in relation to the higher levels of the geographical hierarchy, allowing us to match the wards to a main place. Conveniently, 'Soweto' is one of the defined main places.

In order to map wards to their corresponding main place, spatial joins were constructed using main place shape files and ward shape files. The overlay tool from ESRI (2019) was used, that looked for intersecting main places and wards. The extracted output was then restricted to the Gauteng region. As electoral ward boundaries were not constructed to align with main place boundaries, some wards were

mapped to more than one main place and a matching rule was then constructed based on geographical area, i.e. the ward was assigned to the main place where more than 50 percent of its geographical area is contained. Wards that could not be credibly assigned to a single main place (e.g. split between many main places) were dropped from the analysis.

We note that prior to this, a similar mapping was constructed to administrative regions (where Soweto constituted one of the regions) with the aim of aggregating ward level data to this level. An administrative region is a breakdown separate from the geographical structure mentioned above and has only been defined for the City of Johannesburg (COJ) local municipality, a subset of Gauteng (City of Johannesburg, 2019). There are only seven administrative regions in the City of Johannesburg, hence limiting the donor pool for construction of the counterfactual. One of the requirements in the counterfactual matching is that the values of the variables used to form the weights for the treated unit cannot be outside any linear combination of the values for the donor pool. With Soweto having a larger unemployment rate than any other region in the COJ, we were unable to obtain a 'good' match for the pre-treatment period. Hence main place was used to increase the possibility of finding a suitable match, with a different set of donor pool.

With the ward to main place mappings, QOL and census data were then aggregated to main place level, keeping the variables of interest and ensuring the definitions are the same across the different datasets, for example the definition of 'high school education and above' or household facilities like 'access to a toilet'. A panel dataset was then constructed across the time periods.

Multiple Treatment Units This level of analysis takes each of the 41 wards in Soweto as treatment units and estimates the effect for each, allowing us to identify heterogeneous effects within the township. It also serves the purpose of expanding the donor pool, as the limited number of donors in the main place dataset hinders the quality of the synthetic control. Expanding the donor pool gives us variation in terms of unemployment levels; the 4 mainplaces have larger unemployment levels than Soweto, but there might be individual wards in the groups that have more similar unemployment to Soweto.

In order to be able to run this analysis, we construct a second dataset on the ward level, excluding data from QOL 2009. The exclusion is necessary for the geographical boundaries to be consistent across the dataset. Even though excluding already limited data from our panel is not ideal, this level of disaggregation allows us to zoom in and distinguish between smaller geographical units.

Our primary motivation in breaking down the analysis into smaller dimensions stems from the fact that the wards that make up Soweto are different in their access to treatment (due to size, shape and overall distance to the Rea Vaya stops). Therefore, we acknowledge that Rea Vaya will be a weaker substitute to existing alternatives in some areas and expect the take-up rate to differ across wards. Following from this,

after estimating the individual treatment effect for each treated ward, we first want to understand whether the treatment effect has any relationship with treatment intensity, i.e distance to nearest BRT stop.

In order to assign treatment intensities to wards, we construct a distance to BRT measure. Using ESRI (2019), we first find the centroid of each ward, which in some cases lie outside the ward, and measure the distance from the ward centroids to the nearest BRT stop. As an example, Figure 25 shows the BRT route and stops, along with ward centroids that are 0.7 kilometers away from the nearest stop. We then split the wards into three roughly equally-sized groups; Group 1, those that are at most 0.5 kilometers away from BRT, Group 2, those that are 0.5 to 1 kilometer away, and Group 3, those that are further away than 1 kilometer (See Table 11). We plot the estimated effects against treatment intensities, and we observe a negative relationship with distance and post-treatment RMSPE showing a 40 percent negative correlation. Figure 3 forms the basis of our motivation to pool estimates in groups. We pool the estimates by taking mean averages across these three treatment groups as well as the entire treatment pool.

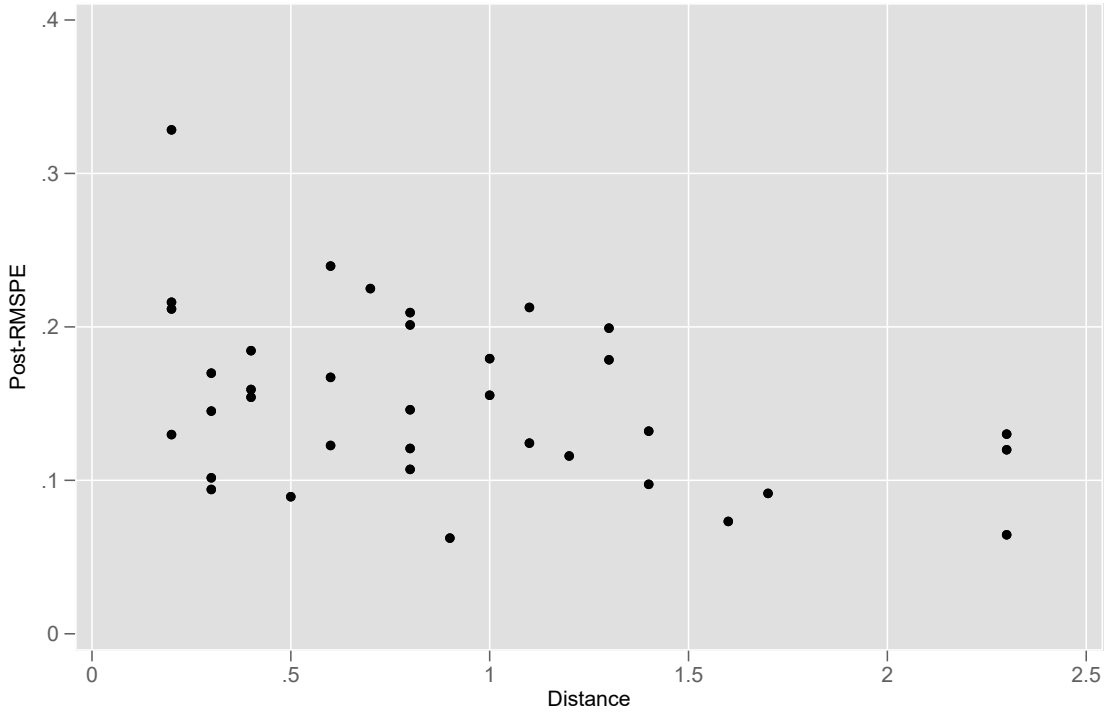


Figure 3: Treatment intensity vs effect size

Kernel Balancing Approach We also provide basic kernel balancing estimates in an attempt to mitigate the interpolation bias arising from the limited pre-treatment data we have. Since we have very few pre-treatment periods, we have the possibility of imposing more constraints on our optimization problem without compromising on the matching quality. Moreover, the methodology optimizes the number of principal factors

the matching will be based on, leaving small room for the researcher to manipulate the process. Kernel balancing assumes the predictor variables to be time-invariant, therefore a modified version of the ward-level panel was constructed to replace the covariate values with their mean across the years. We note that this is consistent with our original methodology, since we have only considered average values as predictors for the SCM.

Outcome Variable In line with our research question, we select unemployment rate and female unemployment rate as outcome variables. Unemployment rate is defined as the number of individuals that are actively seeking employment (in the 4 weeks prior to the date of the QOL survey or census) as a fraction of the workforce. The workforce is defined as individuals who are employed in the formal or informal sector, own a business, or are unemployed. This definition extends easily to female unemployment rate.

Predictors We select the following predictor variables for our analysis:

- Lagged unemployment rate
- Percentage of the population that have completed at least high school education (standard 12)
- Percentage of white or Indian/Asian households
- Percentage of female-headed households
- Household development index

In the case of estimating the effect on female unemployment, female unemployment rate and the corresponding educational indicator are used.

These predictor variables are chosen as they are believed to have large explanatory power with respect to unemployment levels. Variables that could not be standardized across the two different data sources were not considered for the analysis. Even though this limits our choice of predictors, the final list of indicators are in line with the previous labor market in South Africa literature (Dinkelman, 2011).

- It is common to include the lagged values of the outcome variables in the predictor set as lagged values are most likely to capture the factors influencing post-treatment outcomes. However, Kaul, Klößner, Pfeifer, and Schieler (2015) note that using all available values of lagged outcome variable in the synthetic control will render the use of other covariates unnecessary. We note that this is not relevant in our case, as we only consider aggregated pre-treatment values in all specifications. This allows us to include other predictor variables.

- The education rate represents the percentage of ward population with at least grade 12 level education, also known as 'Matric'. This indicator is chosen because the effects of educational attainment on probability of finding employment is the strongest post-matric (Bhorat, Hirsch, Kanbur, and Ncube, 2014; Mlatsheni and Leibbrandt, 2014). However, we note that this is by no means a perfect indicator of whether the student has acquired the literacy and numerical skills required in the labor market (Graham, 2007).

- Percentage of white/Indian/Asian households, clearly, capture the racial composition of the wards. This is included since racial disadvantages are very pronounced in the South African labor market.

- Percentage of female-headed households were included as these households are at greater risk of poverty. In addition to this vulnerability, women bearing the sole responsibility for the household might have implications for the female labor participation rate (Kishor and Neitzel, 1996).

- The household development index was constructed using an inverse covariance weighting method, and takes the following binary inputs – whether the household has a toilet, electricity for lighting, piped water and whether the dwelling type is formal. The index is aimed at aggregating the variables that are assumed to be related to unemployment rate in a similar way, and this method is particularly useful when constructing indices that is believed to represent the same latent factor (Samii, 2019a) This is likely to be the case in the different indicators regarding household dwellings; they are all associated with an overall increase in urbanization and/or income, however might be capturing different pieces of the overall picture. The Stata command `make_index_gr`, developed and made public by Cyrus Samii (Samii, 2019b), was used for construction of the index score.

Variables excluded on purpose are; age cohort and income, as we believe there is too much confounding factors at work for these indicators to be informative in constructing a counterfactual to the treatment unit.

Donors A necessary condition for the synthetic control method is the 'no interference' condition (Abadie et al., 2010), that requires that no unit in the pool of potential donor units have a similar intervention and that the intervention in the treated unit cannot affect the outcome in the pool of donor units. In order to ensure that this condition is met, a spatial join was constructed using the overlay tool in ESRI (2019) to identify all wards in Gauteng (excluding wards in Soweto) that have had a similar intervention over the relevant time period, within a 3 kilometer radius.

Various initiatives by governing bodies related to transport infrastructure improvement as well as transport interventions in the past 10 years are covered in detail in the Strategic Integrated Transport Plan of City of Johannesburg 2013 (Johannesburg Strategic ITP Framework, 2013). We use this document to

identify the other significant interventions in CoJ, as we need to create a donor pool of never-treated areas.

- Metrorail: The Johannesburg Strategic ITP Framework (2013) reports that Johannesburg's commuter rail network, Metrorail, has long been known to provide poor, unsafe and unreliable service; it does not fit entirely with present day residential and economic nodes and journey times are not competitive with other modes (Johannesburg Strategic ITP Framework, 2013). In line with this, the need to upgrade and improve the existing system has been acknowledged and several plans to do so were drafted, however we were not able to find evidence that the plans were followed through. The interventions we are aware of are minor, and we do not consider them to constitute a significant effect such that we would need to exclude these wards from our analysis.

- Gautrain: Gautrain is a modern, high-quality and relatively more expensive, high-speed rail connection linking Johannesburg and Tshwane provinces and Tambo International Airport. The roll-out was done from June 2010 to June 2012. Considering the scale and implementation year of this network, we exclude the wards that the Gautrain line goes through from our analysis.

- Johannesburg has a widespread network of bus services called Metrobus. Even though an Integrated Rapid Public Transport Network (IRPTN) plan was developed for upgrades to the whole network for each municipal area in Gauteng, Johannesburg Strategic ITP Framework (2013) reports that none of these plans have been implemented. However, Metrobus introduced some new services independent of the plan (Johannesburg Strategic ITP Framework, 2013). We exclude the wards with a Metrobus stop as well from our set of untreated wards.

The significant public transport interventions are identified as Metrobus, Gautrain and parts of the BRT line that do not fall in Soweto; the 'treated' wards were then dropped from the original dataset. We further restrict the dataset to wards in City of Johannesburg, due to other transport interventions being very difficult to identify, so our final donor pool is made up of 15 wards (corresponding to 4 main places) within City of Johannesburg that are more than 3 kilometers away from a Gautrain, BRT or Metrobus access point (bus stop or train station).

Selection of predictor and donor weights In the literature, there is currently no consensus on what constitutes a good fit, or similarity between the treated unit and synthetic control. One common tool is choosing the specification that minimizes the pre-treatment RMSPE, as it indicates the best pre-treatment fit possible. Identifying multiple specifications and reporting the associated pre-RMSPE for each, as well as the weight of predictors for each specification chosen allows us to make a data-driven choice. In the single treatment unit analysis, all combinations of predictor variables are tested using the synthetic control method, and the specification that results in the smallest root mean square prediction error in the pre-treatment

period is identified. Even though it is possible to include any linear combination of predictor and outcome variables for the main specification, we only consider the averaged values across all pre-treatment years to overcome the noise that is generated from the sparseness of our data. The chosen specification is then used in all further analysis with 'overall unemployment' as the outcome variable. We repeat this selection process to choose the specification for female unemployment rate analysis, as it is possible that the determinants of female labor force participation are captured by different predictor variables.

Software Two Stata packages, `synth` and `synth_runner`, along with implementation code by Cunningham (2018) were used for the analysis. The former was developed by Abadie et al. (2010) and released with the original SCM paper (looking at the effects of California's ban on smoking), the latter is a wrapper for `synth` developed by Galiani and Quistorff (2016). Both construct counterfactual outcomes by selecting weights of predictor variables and donors, reports RMSPE and balance, reports estimated effects and p-values for statistical inference. The R package `kbal`, developed by Hazlett and Xu (2018), was used for the kernel balancing method.

6 Results

6.1 Single Treatment Unit

We first run the synthetic control method with the unemployment variable using all possible specifications to assess the fit. We then choose the specification with the smallest RMSPE and proceed to run the inference with the preferred specification. In this case, the preferred specification consists of the following predictor variables: lagged unemployment rate, percentage of female-headed households and percentage of Indian/Asian/white households, and has a RMSPE value of 0.0324357.

In Figure 4, the graph on the top left shows the raw unemployment rate for the treated and donor units, while the graph on the top right shows the unemployment rate for the treated unit and the counterfactual or 'synthetic control'. The graphs on the bottom row show the measured treatment effect for the treated unit as well as the placebos (donor units). All the effects reported are Intention to Treat (ITT) effects, as we do not observe BRT usage. Figure 5 shows the p-values in each post-treatment period. In the post-treatment period, the unemployment rate of the treated unit is below that of the counterfactual after 2012, but greater than the counterfactual between 2010 and 2012.

The treatment effect is positive in the first post-treatment period (2011) and negative in the other three (2013, 2015, 2018). This implies that the BRT had a decreasing effect on unemployment in the second to fourth post-treatment periods. Although the treatment effects on the placebo units vary largely in the

post-treatment period, a visual inspection indicates that they mostly lie above that of the treated unit from 2013 to 2018. However, an inspection of the p-values in Figure 5 reveals that this result is insignificant at the 5 percent level, with p-values ranging from 0.25 to 1.

We note that the panel consists of a compilation of seven different datasets, spanning two main sources (the QOL and census). Hence we standardize the outcome variable to account for systematic effects, as recommended by Ferman and Pinto (2018).

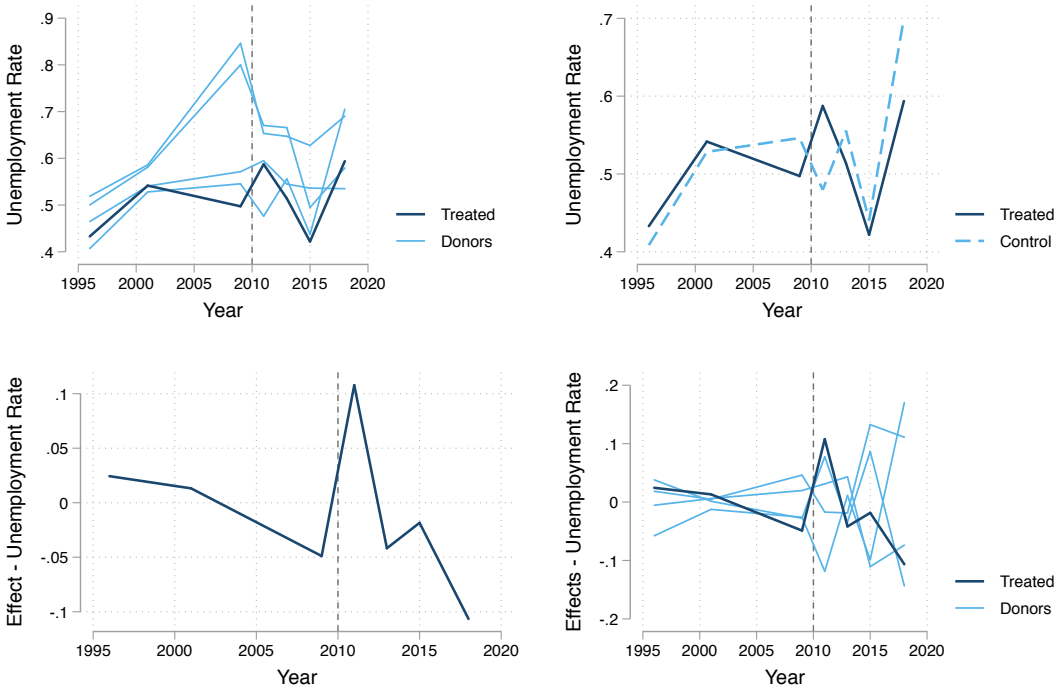


Figure 4: Treatment effects (raw)

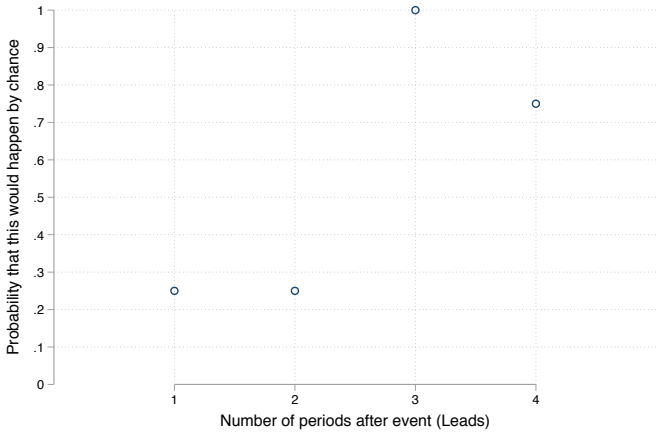


Figure 5: Post-treatment p-values (raw)

The process is then repeated after standardizing the unemployment variable. This is done to address any systematic variation in unemployment rates across the different data sets used. As in the first case, the preferred specification consists of the following predictor variables: lagged unemployment rate, percentage of female-headed households and percentage of Indian/Asian/white households, but has a larger RMSPE value of 0.3279217.

In Figure 6, the graph on the top left shows the raw unemployment rate for the treated and donor units, while the graph on the top right shows the unemployment rate for the treated unit and the counterfactual or 'synthetic control'. The graphs on the bottom row show the measured treatment effect for the treated unit as well as the placebos (donor units). Figure 7 shows the p-values in each post-treatment period. Again, the treatment effect is positive in the first post-treatment period (2011) and negative in the other three (2013, 2015, 2018) and a visual inspection of the treatment effect on the placebo units indicates that they mostly lie above that of the treated unit from 2013 to 2018. However, an inspection of the p-values in Figure 7 reveals that this result is insignificant at the 5 percent level, with p-values ranging from 0.25 to 1. While the scale of the effects in this case differ from the previous case, the relative differences are negligible.

As mentioned before, the inference method we used is limited in its preciseness and interpretation; the p-values reported in Figure7 are merely indicative of relative rank of the treatment effect in the distribution of placebo effects. In our case, having only 4 donor units for 1 treatment unit means that there are only 5 possible values that the p-value can take on. At conventional levels, only the biggest treatment effect will be statistically significant. This limitation is important to keep in mind while interpreting the significance of the treatment effects, and is a motivating factor for us to use disaggregated geographical data.

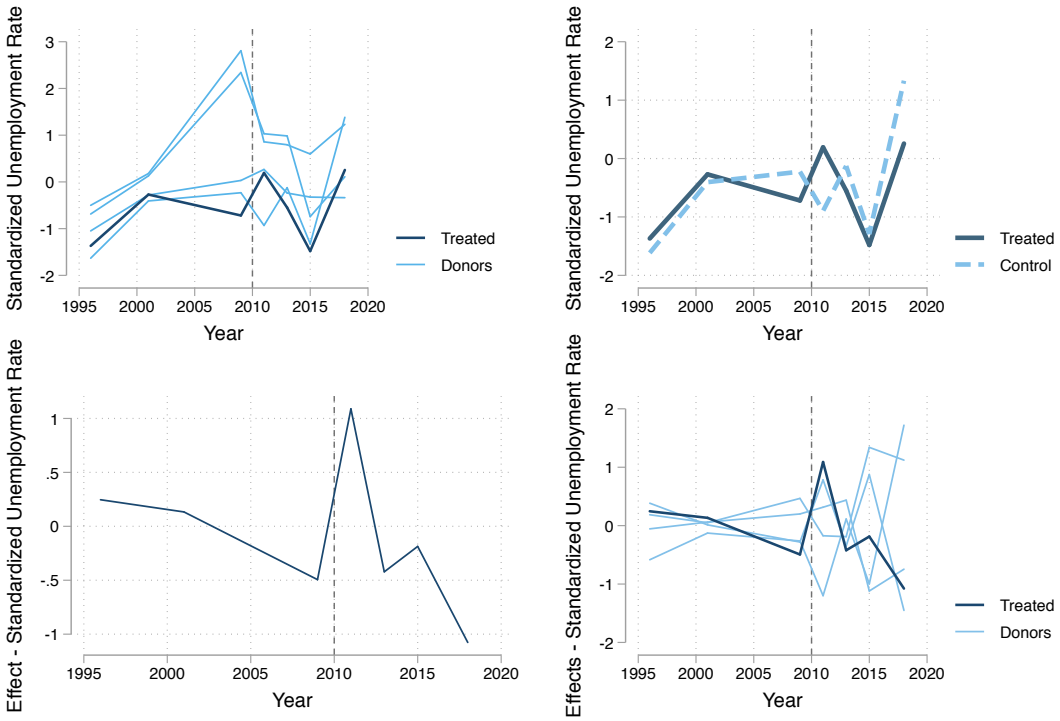


Figure 6: Treatment effects

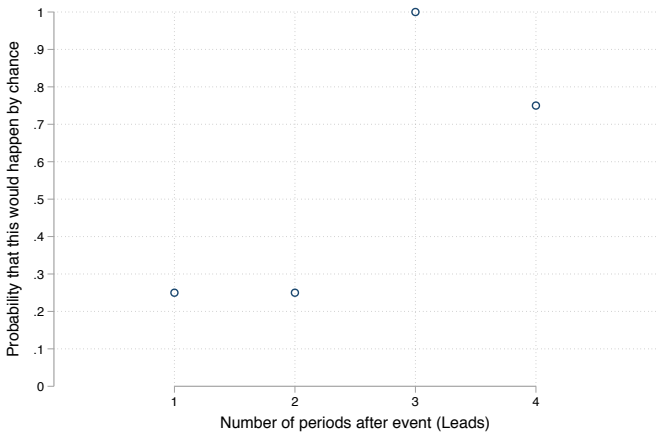


Figure 7: Post-treatment p-values

Table 1 shows values of the predictor variables for the treated and counterfactual units, while Table 2 shows the donor weights chosen for each unit in the donor pool. We note that only units 1 and 2 are used to construct the counterfactual. This is in line with what is observed in Figure 6, where the unemployment rate of these donor units 1 and 2 are closest to that of the treated unit, while the remaining two units have much larger unemployment rates in the pre-treatment period.

Predictor	Treated	Synthetic
Std. unemp.	-.7858145	-.7477493
Female-headed HH	.4250081	.359324
Fraction of white	.0016811	.0001836

Table 1: Predictor balance

Donor Units	Weight
1	.029
2	.971
4	0
5	0

Table 2: Donor weights

We run a similar procedure for normalized female unemployment, this time adding two additional predictor variables to the list of possible predictors – the percentage of males with high school education and the percentage of females with high school education. We find a variety of specifications with the same minimum pre-treatment RMSPE value of 1.017585. We select the same predictors for consistency. The results are shown in a similar manner in Figure 8, with p-values in Figure 9. We observe that the effect on female unemployment is negative in all post-treatment periods except the fourth period, and the match quality is visibly lower. This is corroborated by the fact that the pre-treatment RMSPE value is approximately three times larger than in previous case. As in the previous case, with p-values ranging from 0.25 to 1.00, we see that the effect of the BRT on female unemployment is insignificant at the 5 percent level in all post-treatment periods.

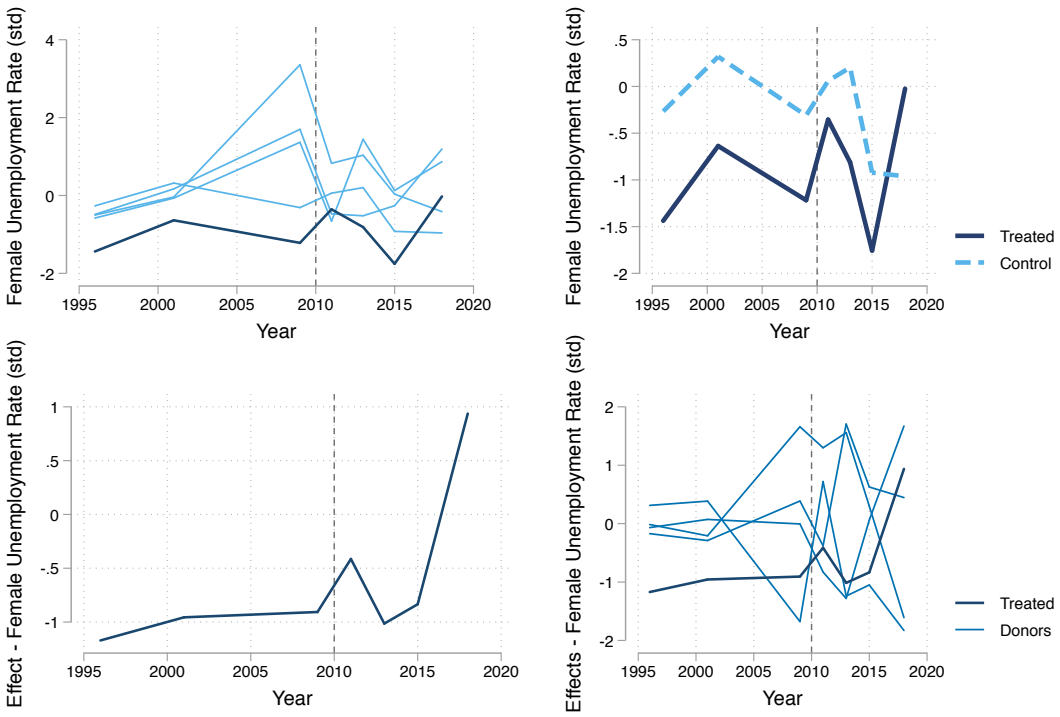


Figure 8: Treatment effects, female unemployment

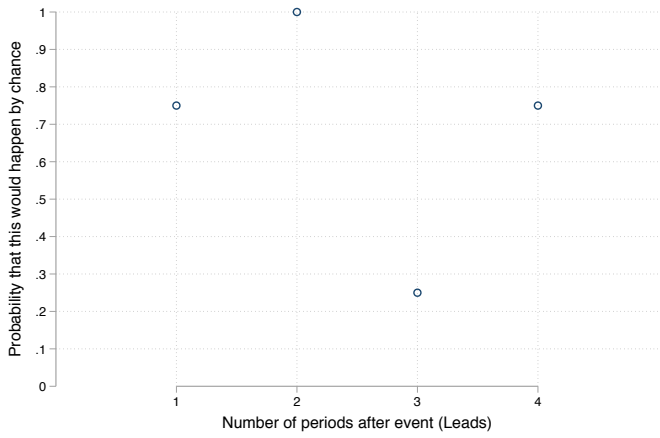


Figure 9: Post-treatment p-values, female unemployment

Table 3 shows values of the predictor variables for the treated and counterfactual units, while Table 4 shows the donor weights chosen for each unit in the donor pool. We note that only unit 1 is used to construct the counterfactual. This is in line with what we observe visually in Figure 8, where three of the four donor units have a relatively larger gap in female unemployment rate as compared to the treated unit.

Predictor	Treated	Synthetic
Std. female unemp.	-.7858145	-.7477493
Female-headed HH	.4250081	.359324
Fraction of white	.0016811	.0001836

Table 3: Predictor balance for female unemployment

Donor Units	Weight
1	1
2	0
4	0
5	0

Table 4: Donor weights for female unemployment

Due to the lack of publicly available data on public transport sources and transport interventions outside of the COJ, the donor pool was limited to the COJ, within which more than half of the electoral wards outside of Soweto were dropped as there was evidence of alternative forms of public transport in these wards. This limited the variation in the values of the predictor and outcome variables across the donor pool. In the single treatment unit case, only four main places were available for construction of the counterfactual, with unemployment rates that were greater than Soweto. This compromised the quality of the matching, which led to high pre-treatment RMSPE values in both unemployment and female unemployment. In light of this, we proceed with multiple treatment units approach to mitigate this issue.

6.2 Multiple Treatment Units

Here we present the results of the SCM with disaggregated data. We start by examining Figures 19 and 20. These figures present the raw data for our two outcome variables, and contrasting this to the single treatment approach, we see that the similarity of the trends and absolute values for both outcome variables will allow us to approximate treated units' trajectory with a linear combination of donor wards much better. As such, our reported RMSPE for pre-treatment periods will improve in comparison to the single unit approach. Indeed, we are able to approximate the pre-treatment trajectory of each treated unit except one with an RMSPE smaller than 0.05, for both outcome variables. Similarly, Figures 21 and 22 show the values for two predictor variables in the treated wards, fraction of white/Indian/Asian households and female headed households are at reasonably close levels in treatment and donor groups. This allows us to approximate the counterfactual Soweto wards as a linear combination of available donor units. At a maximum RMSPE of 0.03, we are able to match the donor units well in the pre-treatment period as well.[2] These matches are available for visual

[2]We reiterate that there is no consensus on the literature of what constitutes a benchmark for a "good" match, and the RMSPE values are relative to the scale of outcome variable.

inspection in Figures 10 and 13, where we see that the treated and counterfactual groups follow each other closely pre-treatment but differ post-treatment. The differing trajectories for each treatment-counterfactual pair during the post-treatment periods is equal to the estimate of treatment effect.

Figure 11 shows the average pooled treatment effects for three groups as well as placebo effects for each donor unit. Re-iterating that treatment effect is calculated as the difference between the observed and the constructed counterfactual unemployment levels, this figure encapsulates the gist of the methodology in an intuitive manner. An initial glance shows that there is a sharp drop within the first two post-treatment periods for our treatment groups, with effects for treatment group 2 being especially strong at -0.2. All in all, there is a clear downward trend for all treatment groups. However, this is following the general trend, with many placebo units experiencing a similar drop as well, pointing to insignificant results. Before looking at effect sizes in detail by year, we first examine Table 5, which reports the average treatment effect across the years 2011-2018 together with the associated test statistic. For all of our treatment groups, pre-treatment effects are close to 0, as the quality of match is high. While the average treatment effect on unemployment is negative in all groups, we see with p-values well above the conventional levels, only Group 1 has an average treatment effects are statistically significant. Table 12 reports the pooled treatment effects by groups and placebo effects for each donor, for each year. This is to be complemented by Figure 12, which shows the associated p-value for each year. The probability that these treatment effects would be observed purely by chance for each year varies across years, at times as high as 0.5. The interpretation of very high p-values are simply that the pooled post-RMSPE values for that group was among the smallest throughout in that particular year. We observe that even though treatment groups did not have a statistically significant effect in most of the years, the effects for Groups 1 and 2 are significant at 10 per cent level in year 2015.

Table 14 (Appendix) shows the average treatment effect on each of the 41 treated wards. We see that at the 10 percent significance level, 15 wards achieve statistical significance. Of these wards, five belongs to treatment group 1, four belong to treatment group 2, and the last six with statistically significant effect belongs to treatment group 3. These effect sizes vary from 5.3 percent to 17 percent. However, we prefer to give a larger weight to our pooled estimates, as the wards with significant treatment effects seem to share no common characteristics and are of varying treatment intensities, making it difficult to hypothesize regarding the mechanism involved.

Unit	ITT	P-value
Group 1	-.09741002	.0666667
Group 2	-.06957777	.2
Group 3	-.06527806	.2666667

Table 5: ITT and significance for pooled estimates

Turning to female unemployment, Figure 14 shows the treatment effects for female unemployment by group. Firstly, we see that the matching quality has improved significantly in comparison to the single

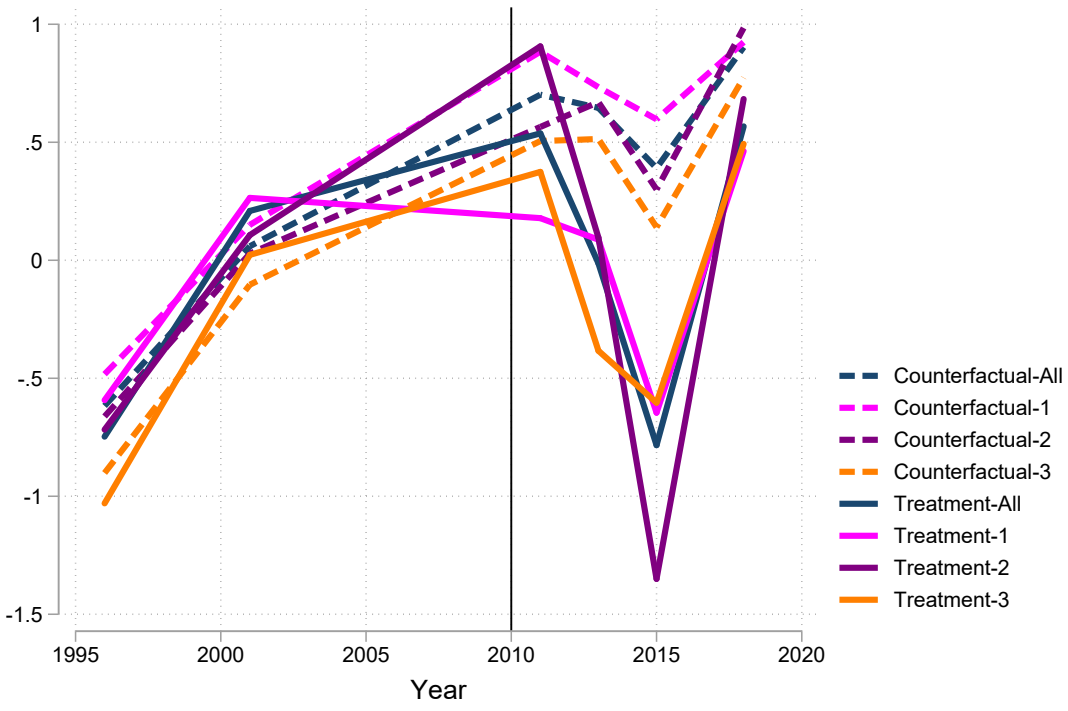


Figure 10: Treatment and counterfactual groups for unemployment



Figure 11: Treatment effects for all units

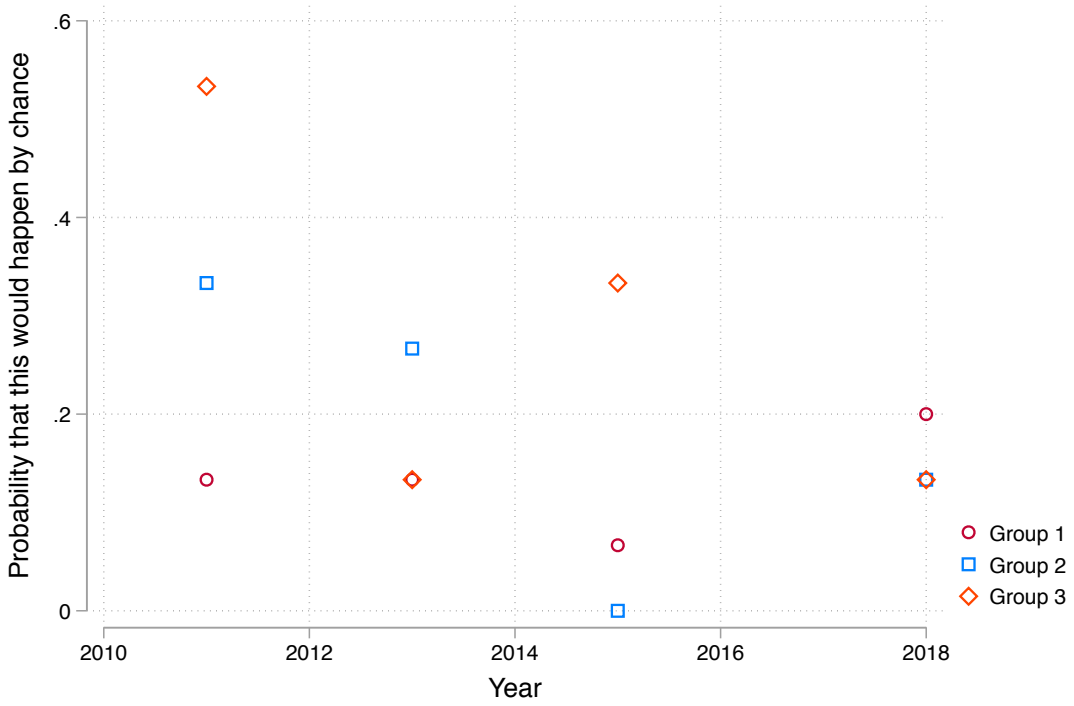


Figure 12: P-values for each year, by group

treatment unit approach. The post-treatment gaps between the actual and the counterfactual are more pronounced compared to overall unemployment; figure even indicates a reversal of trend. However, the absolute differences are smaller, consequently, we observe smaller effect sizes in Table 6. Although Figure 14 shows that these effect sizes lie at the lowest point with respect to the untreated wards, these effects do not seem to differ from 0 in the last treatment period, in fact they could even indicate a positive effect. Indeed we see in Figure 9 that none of the pooled effects are statistically significant. Table 13 and Figure 15 together present the treatment effects and significance by year, and similar to overall unemployment, for only Group 2 in 2015 and Group 3 in 2013 we see the treatment effects are significant. We also report the treatment effects and P-values for each ward in Table 15; nine wards individually have statistically significant treatment effects. However for the same reasons mentioned above we prefer to base our main conclusions on pooled estimates, that is the roll-out of Rea Vaya did not have significant effects on unemployment in Soweto. We further note that precision of the p-values are significantly limited due to the small number of units (15 donor and 1 treatment in each inference set), and a larger number of units will have given more informative p-values.

Unit	ITT	P-value
Group 1	-.06484912	.2
Group 2	-.04587615	.133333
Group 3	-.04351328	.2

Table 6: ITT and significance for female pooled estimates

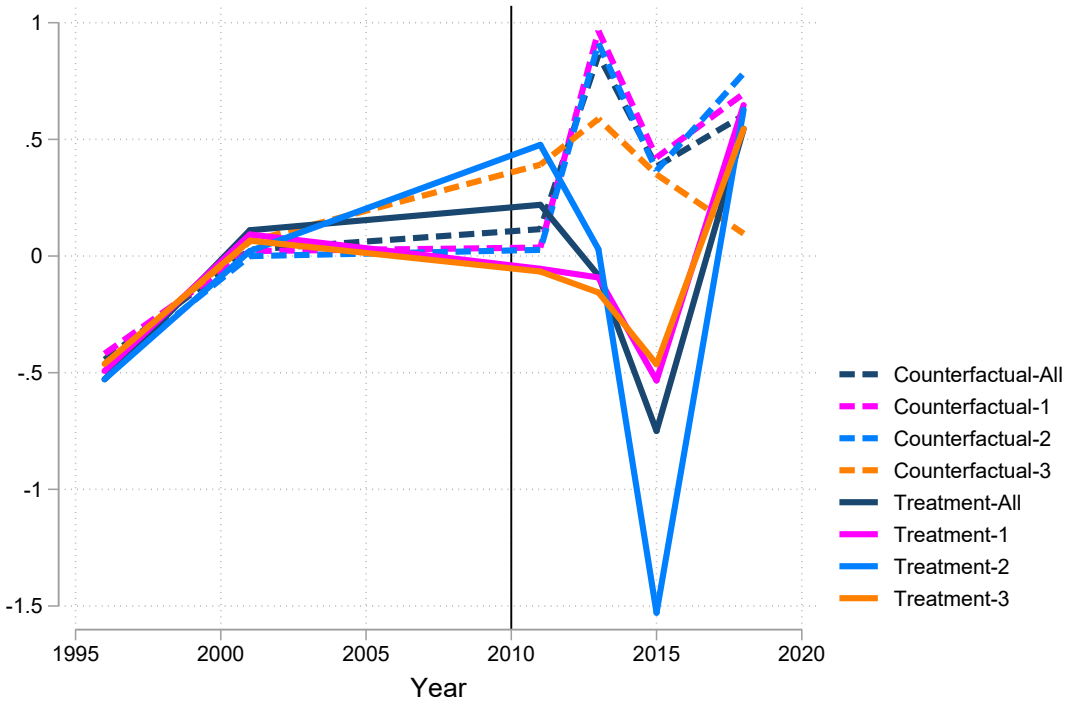


Figure 13: Treatment and counterfactual groups for female unemployment

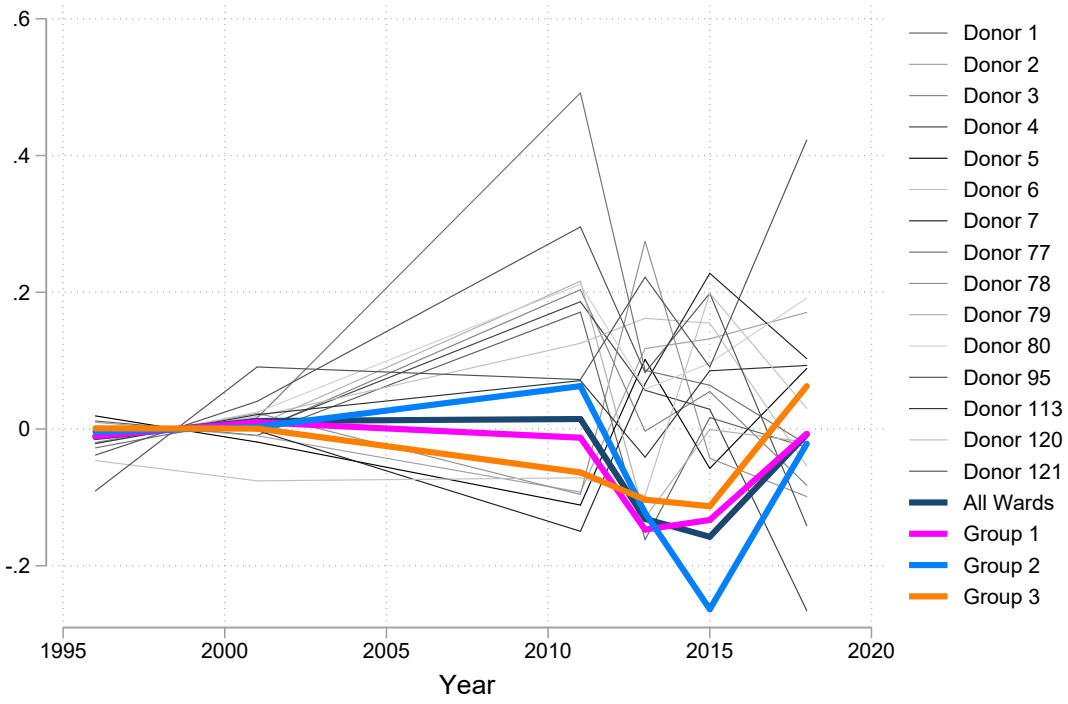


Figure 14: Treatment effects for all units, for female unemployment

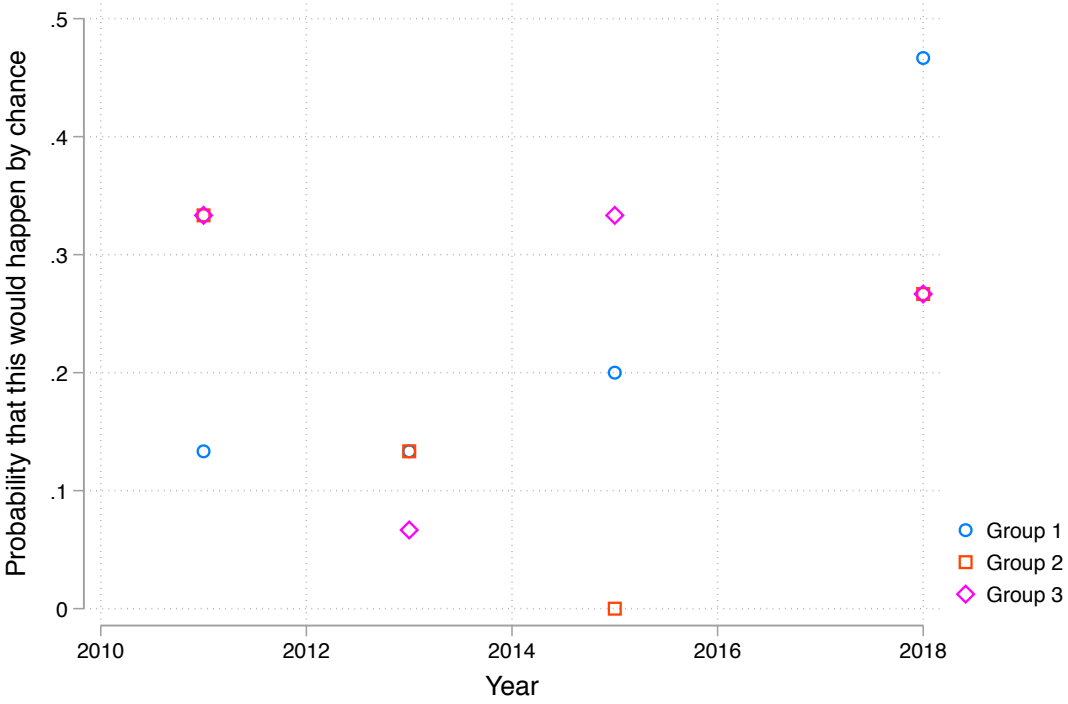


Figure 15: P-values for female unemployment each year, by group

6.3 Kernel Balancing

In this section, we present the results obtained by kernel balancing. Recall that this method seeks balance on higher orders, even if it means a trade-off on the mean balance between treatment and synthetic units. This can be seen in Figure 16, where the differences in means of treatment and weighted control units are actually larger than the difference in means of treatment and un-weighted, raw donor units. The reason we still provide these estimates is due to the fact that kernel balancing matches the general trajectory of the treatment group by imposing more constraints, in our case, more than the constraints imposed by two pre-treatment periods.

The improvement in the pre-treatment match quality is evident in Figure 17, where the control group is denoted by $Y(0)$.[3] Table 7 reports the ITT effects for the pooled estimates of kernel balancing, as well as the year-by-year estimates. For both pre-treatment periods, we see that the treatment effect is not significantly different from 0. In 2013 and 2015, the ITT effect is estimated at 3 percent and 13 percent respectively, and significant at the 1 percent level. While the effect sizes are comparable to our previous estimates, the significance of the results could possibly be attributed to the higher matching quality for both the donor and treated units generated by kernel balancing. The average ITT for all treated wards across the years is less than 1 percentage point, however statistically significant as well.

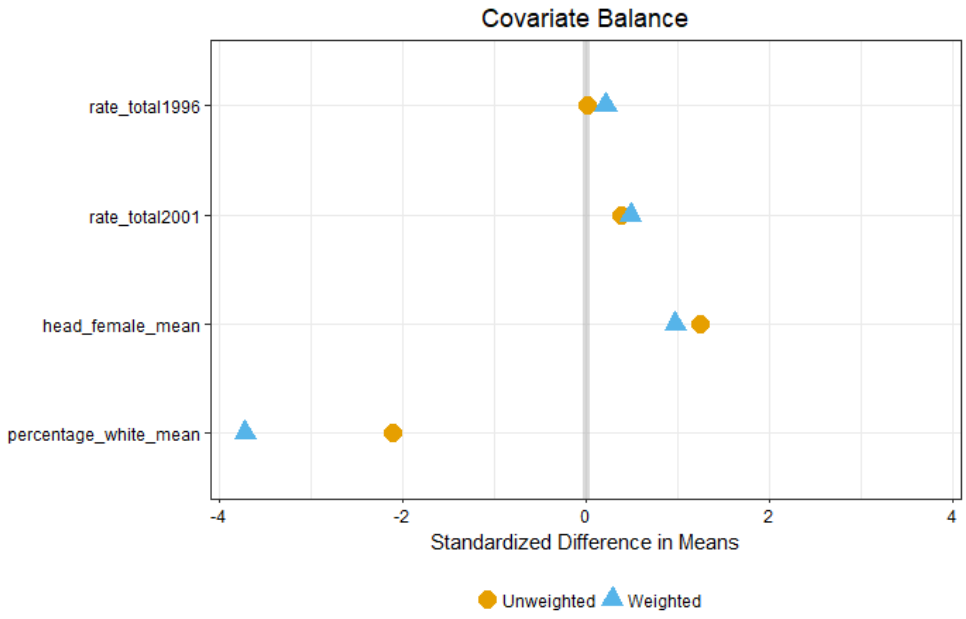


Figure 16: Covariate balance

[3]With the removal of 2009 data due to inconsistencies in ward boundaries, treatment baseline becomes 2001 when visually displayed.

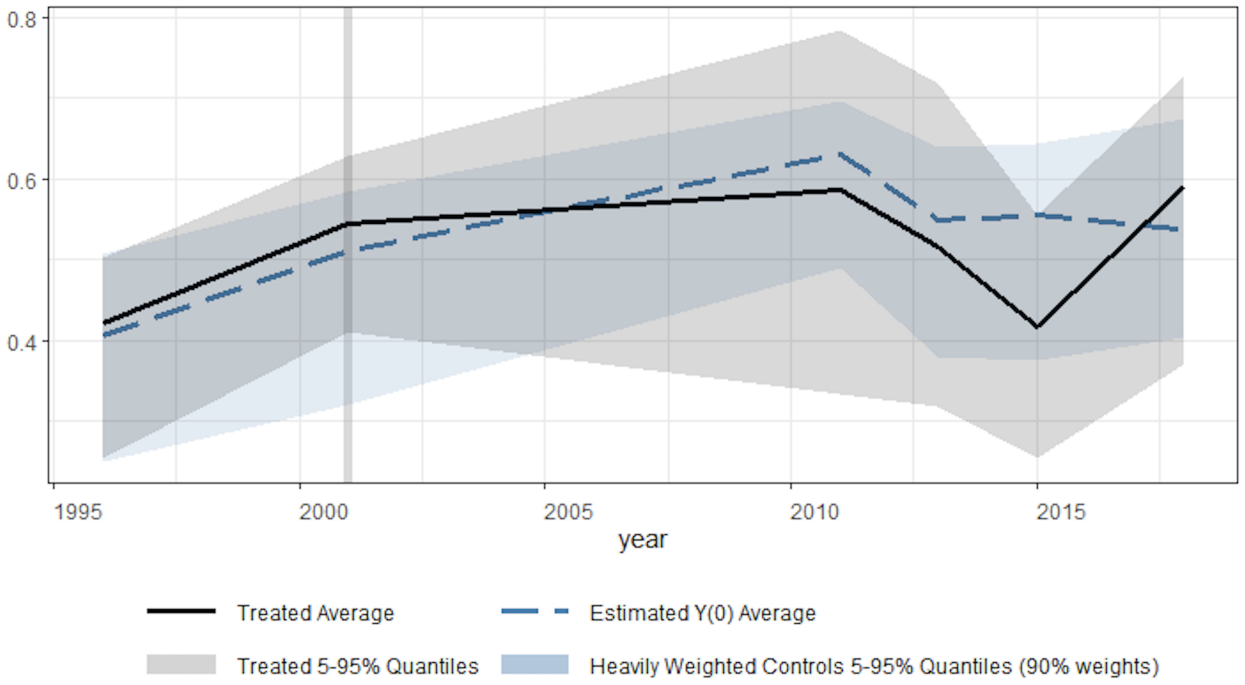


Figure 17: Treated and counterfactual averages

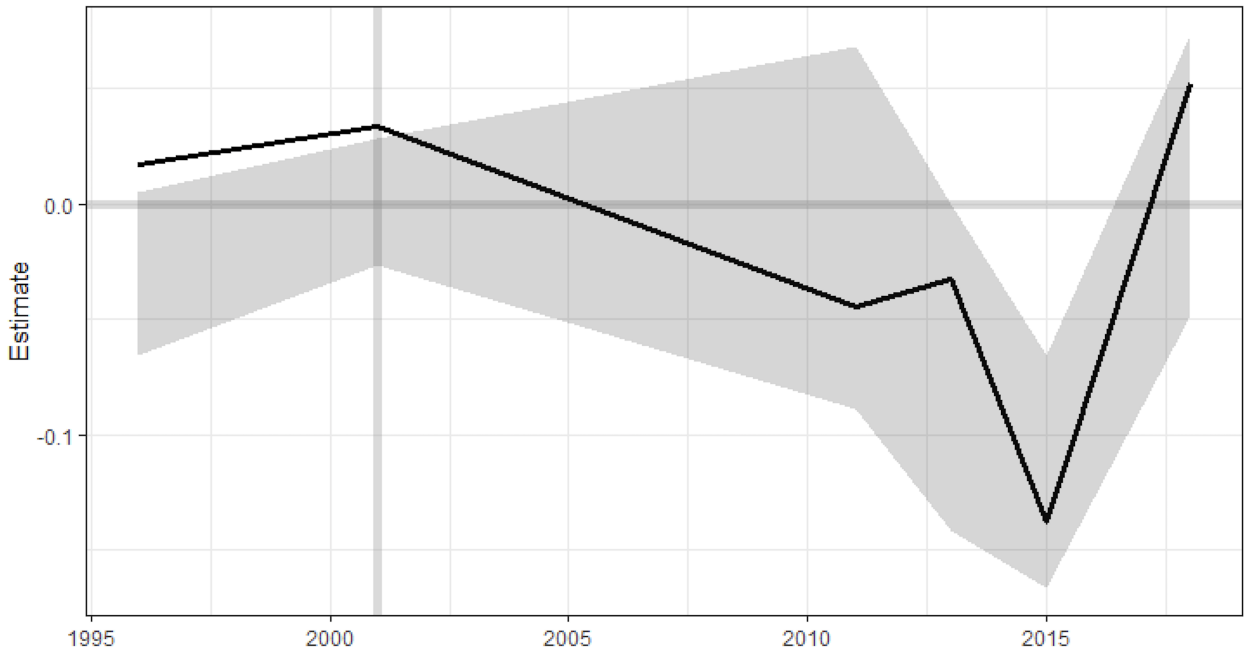


Figure 18: ITT for treated wards

	ITT	S.E.	CI-lower	CI-upper	P-value
Average	-0.04082	0.01669	-0.08566	-0.01676	0
1996	0.01684	0.01765	-0.06542	0.0046291	0.07
2001	0.03335	0.01370	-0.02676	0.0285834	0.86
2011	-0.04457	0.04048	-0.08934	0.0679320	0.53
2013	-0.03238	0.03784	-0.14208	-0.0001188	0.05
2015	-0.13831	0.02623	-0.16660	-0.0654659	0.00
2018	0.05197	0.03085	-0.04855	0.0731672	0.67

Table 7: Average ITT effect for treatment wards

7 Discussion

Our results have mainly shown that even though we observe large effects across groups and years, we are unable to assert that these effects are not there purely by chance. On the other hand, for a number of individual wards and some groups, we observe significant effects; however they seem to follow no explanatory pattern. Our inference procedure further limits our precision. We proceed with a discussion on the observed results and limitations of the paper, and suggest areas where future research in this field should set their focus on.

7.1 Limitations

As mentioned previously, our methodology and data constraints prevent us from disentangling the mechanisms through which access to jobs is affected as well as capturing the first-stage effects on the intervention. We first reiterate the limitations of our methodology, then, bearing these limitations in mind, we discuss possible reasons that (i) Rea Vaya has not increased job accessibility, therefore had no effect on employment levels and (ii) our methodology hasn't picked up on the effects that Rea Vaya had.

7.1.1 Internal Validity

We rely on QOL surveys for our dataset post-2009. Although the sampling was done to be representative on a ward level along the lines of gender and race, it might not be as reflective of the actual unemployment rates on the ward-level as census data is. To overcome this and other potential systematic differences, we standardize our datasets.

Perhaps the biggest limitation in our analysis is the limited number of pre-treatment data available to us, combined with the sparseness of the periods in which the data is available. This limits the information

we can incorporate in constructing the synthetic control, and therefore limits the ability of the SCM to mimic the trajectory of unobserved confounders.

In order to be conservative regarding the treatment status of donor wards, we have limited our dataset to City of Johannesburg only; however this has left us with 15 donor wards. Even though we were able to match the pre-treatment trajectories of Soweto wards reasonably well with these donor units, our inference procedure has been severely limited. The small number of donor wards limits the number of placebo effects, which are used to construct the distribution and consequently the rankings. Increasing the number of donor wards by confirming that these wards have not received any treatment throughout the time frame of the dataset could improve our results.

The need to match the definitions of variables between the census and QOL survey dataset has limited our choice of variables. The census, as expected, does not collect as detailed information as a household survey does. The limited choice of predictor variables, combined with the limited pre-treatment data to be validated, might have resulted in poor approximation of the counterfactual outcome.

A particular problem in our results is potential interpolation bias, which we specifically attribute to the 2008-2009 Global Financial Crisis. The 2008 crisis had pushed South Africa into a recession, with the already slowing economy losing close to 1 million jobs during the crisis (Selassie, 2011). In addition to this, by 2011 the ratio of employment to the working-age population had declined from around 45 percent in 2008 to 40 percent in 2011, implying that employment had remained below the pre-crisis level (Mlatsheni and Ranchhod, 2017). If Soweto, or donor units that have been heavily weighted in synthetic Soweto, have been disproportionately effected by this crisis, this would mean that the assumption of linearity between predictor and outcome variables does not hold across all units in our datasets. If this is indeed the case, it will result in interpolation bias and will be a limiting factor on the synthetic control to mimic the trajectory of Soweto.

7.1.2 Theoretical Framework

The reach and scale of the Rea Vaya (phase 1A) pales in comparison to that of the minibus taxis that densely populate the Soweto township and provide greater accessibility than the Rea Vaya in terms of the walking distance to the transport mode. The taxi map is shown in Figure 28. It might be the case that the individuals who have actually received treatment (in terms of increase in accessibility in any form) has remained a small portion of the population. This would indicate that even if Rea Vaya has indeed helped a small sample of individuals to get jobs, we would not be able to observe those in our study. We note that the Rea Vaya rollout is still ongoing and the scaling of its operations in later phases might lead to greater effects, however this is yet to be seen.

The key mechanism through which decreased unemployment was expected in this context is the

increased individual rationality of searching for and accepting employment in the Johannesburg city center via a reduction in commuting costs. However, while the Rea Vaya may have driven cost savings in absolute terms, there is evidence that the transport burden remained considerably high when compared to income levels. According to De Lannoy et al. (2018), the average job search cost among a sample of youth was almost double their monthly per capita household income, which was at a level where food security was a reality. In this context, individuals are likely to be well below the job-seeking threshold described in Theory 1. In addition, there are also non-transport costs incurred when searching for work, particularly in this context where access to information is a challenge. According to De Lannoy et al. (2018), these costs include 'payments at internet cafes, printing and postage costs', and the cost of looking for information on the internet is particularly high. As a result, it is likely that the costs involved in searching for a job, expanding the area of search, or commuting costs associated with taking up jobs far away, have remained above the critical threshold for many job-seekers, which would mean that even if there was an increase in accessibility levels for the individuals, it failed to translate to any significant secondary outcomes.

While the aim of this paper is to single out the effect of a transport intervention on unemployment in the South African context, it is important to understand that transport disadvantage is a barrier that lies above other existing barriers that the socially disadvantaged are faced with and exacerbates their plight in seeking for and successfully procuring work. In the theory described previously, the individuals conduct job search taking into account probability of finding suitable jobs. In this context, Soweto residents have poor access to networks and low education levels which negatively affect their access to employment. Consequently, even if the job search costs have gone to levels that would push the residents above the threshold, the perceived probability of finding a job might have stayed as a binding constraint on job search behavior. The joint effect of removing these barriers together with transport barriers is yet to be determined.

The outcome variable used in our analysis, unemployment rate, captures the portion of the workforce that was actively seeking employment. While there was no significant effect in that variable, we were unable to capture the employment effect on the population that were not initially actively seeking work. This group is often referred to as 'discouraged' individuals, which comprises ex job-seekers that have given up on their job search for a variety of reasons. In the South African context, discouraged job seekers are a prominent group, and according to the De Lannoy et al. (2018), 'location and the cost of looking for employment are indeed constraining factors for job search among unemployed, non-searching respondents', with over 70 percent of this group citing spatial constraints as a barrier to searching for employment in the 2005 Labor Force Survey (De Lannoy et al., 2018). In addition to any uncaptured effects on discouraged workers landing employment, in the event of their entrance into the workforce and unsuccessful job search, there would be an upward bias in the measured unemployment rate. Due to a discrepancy across datasets in how discouraged job seekers are defined, we were unable to create a consistent indicator to capture the employment effect on this group.

In addition, the Rea Vaya phase 1a was implemented in stages from late 2009 to 2010, while the post-treatment period in our analysis begins in 2011 as the QOL survey is conducted every two years on average. This implies that any treatment effect during the two-year roll-out period prior to 2011 is not captured in our analysis.

7.1.3 Further Research

We remain convinced that the link between labor market outcomes and transport equity is strong in South Africa, and that further research could immensely benefit policy-making.

Accessibility measures in an African context Transport accessibility and resulting economic outcomes is a fairly novel, under-researched topic. Particularly in the Sub-Saharan Africa context, where transport costs are disproportionately high and many cities are characterized by rapid but fragmented urban growth (Lall, Henderson, and Venables, 2017). In such a context, partialling out the effects of each component (accessibility, affordability, availability and acceptability) on secondary outcomes is a particularly important area to study. Transport subsidies, rapidly gaining traction in policy research, allows the researcher to create and control the first stage of the theoretical framework we explore in this paper and make sure that the job-seeking threshold is crossed. One example of such a forthcoming study is by Banerjee and Sequeira (2019).

Provision of "good" jobs The theories and mechanisms we describe apply not just to widening the pool of potential jobs, but widening the pool of better jobs in terms of social security and formal job training. Measuring whether transport improvements lead to an increase in the quality of jobs held by the disadvantaged, and how this extends to the overall happiness, life satisfaction, inequality and social integration within the South African society will be highly relevant in guiding transport policy.

New integrated fare collection system A system of tap-and-go cards that cardholders can use for travel in all transport systems, as well as for small purchases in retail outlets, is being rolled out in South Africa through an integrated fare collection system (IFCS). This is a safer and more consumer-friendly alternative to the existing system; it limits cardholders' exposure to ZAR 200 per day, and prevents the fares from fluctuating in an informal manner (Woolf and Joubert, 2013). The IFCS has been surprisingly well-received and is gaining traction quickly, with the taxi industry having already been fitted with card readers (Woolf and Joubert, 2013). As such, it would be informative to examine the effect of this technology injection on both the BRT take-up rate as well as on labor market outcomes.

Rea Vaya 1c Phase 1c of the Rea Vaya is expected to be complete by 2030, and the scale and reach is a lot more ambitious than the current system. Collecting relevant individual level data (including travel habits, income, job prospects) prior to this intervention, or repeating the analysis in this paper with Enumeration Area level census data with more frequent data points, will enable future researchers to assess the effects of this upcoming expansion.

8 Conclusion

In this paper, we look at the employment effect of a public transport intervention in Johannesburg, South Africa, in the form of a BRT implementation. We hypothesize that transportation costs hinder the optimal job-searching effort and take up of employment for the low-income, low-skilled Soweto population, and ask whether an increase in the availability and affordability of public transport could improve these adverse labor market outcomes. This question was motivated by the unique history of spatial segregation of the black communities during the apartheid regime, which led to the development of townships which have limited access to employment opportunities in the city center.

We exploit the intervention of the BRT line connecting the township of Soweto, where unemployment is high, to the city center. Using a synthetic control method, we construct a 'synthetic Soweto' in order to examine the effects on the township as a whole as well as incorporate treatment intensities on a smaller, ward-level aggregation in the form of distance to the nearest BRT stop. Although we are able to match the pre-treatment trajectories of the treated variables fairly well, the limited availability of pre-treatment data limits the validity of the post-treatment counterfactual. We therefore explore kernel balancing as a mitigation tool.

We find that there is a drop in unemployment that is correlated with treatment intensity, and while the average effects are not statistically significant at conventional levels, there are significant decreases in unemployment in selected years.

References

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, *105*(490), 493–505.

Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the Synthetic Control Method. *American Journal of Political Science*, *59*(2), 495–510.

Abadie, A. & Gardeazabal, J. (2003). The economic costs of conflict: A case study of the Basque Country. *American Economic Review*, *93*(1), 113–132.

Abel, M. (2017). *Labor Market Discrimination and Sorting: Evidence from South Africa*. Policy Research Working Papers. The World Bank.

Acemoglu, D., Johnson, S., Kermani, A., Kwak, J., & Mitton, T. (2016). The value of connections in turbulent times: Evidence from the United States. *Journal of Financial Economics*, *121*(2), 368–391.

Allen, H. (2013). Africa's First Full Rapid Bus System: The Rea Vaya Bus System in Johannesburg, Republic of South Africa, 17.

Altman, J. & Hinkson, M. (2007). Mobility and Modernity in Arnhem Land: The Social Universe of Kuninjku Trucks. *Journal of Material Culture*, *12*(2), 181–203.

Athey, S. & Imbens, G. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, *31*(2), 3–32.

Banerjee, A. & Sequeira, S. (2019). *Transport subsidies and job matchmaking in South Africa*.

Bank, W. (Ed.). (2011). *Gender equality and development*. World Development Report. OCLC: 762210346. Washington, DC: World Bank.

Ben-Michael, E., Feller, A., & Rothstein, J. (2018). The Augmented Synthetic Control Method. *arXiv:1811.04170 [econ, stat]*. arXiv: 1811.04170.

Bhorat, H. (2012). A Nation in Search of Jobs: Six Possible Policy Suggestions for Employment Creation in South Africa. *SSRN Electronic Journal*.

Bhorat, H., Hirsch, A., Kanbur, S. M. R., & Ncube, M. (Eds.). (2014). *The Oxford companion to the economics of South Africa* (First edition). OCLC: ocn898753209. Oxford, United Kingdom: Oxford University Press.

Bhorat, H., Kanbur, R., & Mayet, N. (2012). Estimating the Causal Effect of Enforcement on Minimum Wage Compliance. *Review of Development Economics*, *16*(4), 608–623.

Blumen, O. (1994). Gender differences in the journey to work. *Urban Geography*, *15*(3), 223–245.

Bocarejo, J. P., Portilla, I., & Meléndez, D. (2016). Social fragmentation as a consequence of implementing a Bus Rapid Transit system in the city of Bogota. *Urban Studies*, *53*(8), 1617–1634.

Borker, G. (2017). Safety First: Perceived Risk of Street Harassment and Educational Choices of Women, 67.

Brueckner, J. K. & Zenou, Y. (2003). Space and Unemployment: The Labor Market Effects of Spatial Mismatch. *Journal of Labor Economics*, *21*(1), 242–262.

BusinessTech. (2015). White vs black unemployment in South Africa. BusinessTech.

Carlsson, M., Reshid, A., & Rooth, D.-O. (2017). Neighborhood signaling effects, commuting time, and employment: Evidence from a field experiment. Linnéuniversitetet.

Cavallo, E., Galiani, S., Noy, I., & Pantano, J. (2012). Catastrophic Natural Disasters and Economic Growth, 35.

Church, A., Frost, M., & Sullivan, K. (2000). Transport and social exclusion in London. *Transport Policy*, *7*(3), 195–205.

City of Johannesburg. (2019). Administrative regions. Retrieved May 11, 2019, from https://www.joburg.org.za/about_/regions/Pages/City-of-Johannesburg-regions.aspx

Clark, G. L. & Whiteman, J. (1983). Why Poor People Do Not Move: Job Search Behavior and Disequilibrium Amongst Local Labor Markets. *Environment and Planning A: Economy and Space*, *15*(1), 85–104.

Cunningham, S. (2018). Causal Inference: The Mixtape, 328.

De Lannoy, A., Graham, L., Patel, L., & Leibbrandt, M. (2018). What drives youth unemployment and what interventions help?

Dercon, S. & Hoddinott, J. (2005). Livelihoods, Growth, and Links to Market Towns in 15 Ethiopian Villages. *SSRN Electronic Journal*.

Dinkelman, T. (2011). The Effects of Rural Electrification on Employment: New Evidence from South Africa. *American Economic Review*, *101*(7), 3078–3108.

Donaldson, D. (2018). Railroads of the Raj: Estimating the Impact of Transportation Infrastructure. *American Economic Review*, *108*(4-5), 899–934.

Donaldson, D. & Hornbeck, R. (2016). Railroads and American economic growth: A "market access" approach. *The Quarterly Journal of Economics*, *131*(2), 799–858.

Doudchenko, N. & Imbens, G. (2016). *Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis* (tech. rep. No. w22791). National Bureau of Economic Research. Cambridge, MA.

Dube, A. & Zipperer, B. (2015). Pooling Multiple Case Studies Using Synthetic Controls: An Application to Minimum Wage Policies, 60.

ESRI. (2019, May 11). Arcgis online (Version 7.1).

Fagnani, J. (1987). Daily commuting time: The stakes for working mothers in France. *Transportation Research Record*, *1135*, 26–30.

Fanning Madden, J. (1981). Why women work closer to home. *Urban Studies*, *18*(2), 181–194.

Ferman, B. & Pinto, C. (2018). Synthetic Controls with Imperfect Pre-Treatment Fit, 45.

Findley, L. & Ogbu, L. (2011). South Africa: From Township to Town. *South Africa after Apartheid*, 9.

Firpo, S. & Possebom, V. (2016). Synthetic control estimator: A generalized inference procedure and confidence sets. *Sao Paulo School of Economics, EESP-FGV*.

Galiani, S. & Quistorff, B. (2016). The synth_runner Package: Utilities to Automate Synthetic Control Estimation Using synth, 16.

Gauteng City-Region Observatory. (2019a). Quality of life survey i (2009). Retrieved May 11, 2019, from http://gcro.ac.za/research/project/detail/quality-of-life-survey-i-2009/

Gauteng City-Region Observatory. (2019b). Quality of life survey ii (2011). Retrieved May 11, 2019, from http://gcro.ac.za/research/project/detail/quality-of-life-survey-ii-2011/

Gauteng City-Region Observatory. (2019c). Quality of life survey iii (2013). Retrieved May 11, 2019, from http://gcro.ac.za/research/project/detail/quality-of-life-survey-iii-2013/

Gauteng City-Region Observatory. (2019d). Quality of life survey iv (2015/16). Retrieved May 11, 2019, from http://www.gcro.ac.za/research/project/detail/quality-of-life-survey-iv-201516/

Gauteng City-Region Observatory. (2019e). Quality of life survey v (2017/18). Retrieved May 11, 2019, from http://gcro.ac.za/research/project/detail/quality-of-life-survey-v-201718/

Geurs, K. T., van Eck, J. R. et al. (2003). Evaluation of accessibility impacts of land-use scenarios: The implications of job competition, land-use, and infrastructure developments for the Netherlands. *Environment and Planning B*, *30*(1), 69–88.

Gilbert, A. (2008). Bus Rapid Transit: Is Transmilenio a Miracle Cure? *Transport Reviews*, *28*(4), 439–467.

Glaeser, E. L., Kahn, M. E., & Rappaport, J. (2008). Why do the poor live in cities? The role of public transportation. *Journal of Urban Economics*, *63*(1), 1–24.

Gobillon, L. & Magnac, T. (2013). Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls, 56.

Gobillon, L., Selod, H., & Zenou, Y. (2007). The Mechanisms of Spatial Mismatch. *Urban Studies*, *44*(12), 2401–2427.

Gomide, A., Leite, S., & Rebelo, J. (2005). Public transport and urban poverty: A synthetic index of adequate service. *Documento de Trabajo*, (33901).

Gordon, P., Kumar, A., & Richardson, H. W. (1989). Gender differences in metropolitan travel behaviour. *Regional Studies*, *23*(6), 499–510.

Graham, D. J. (2007). Agglomeration, Productivity and Transport Investment. *Journal of Transport Economics and Policy*, *41*, 27.

Grogger, J. & Hanson, G. H. (2011). Income maximization and the selection and sorting of international migrants. *Journal of Development Economics*, *95*(1), 42–57.

Gutiérrez-i-Puigarnau, E. & van Ommeren, J. N. (2010). Labour supply and commuting. *Journal of Urban Economics*, *68*(1), 82–89.

Gyimesi, K., Vincent, C., & Lamba, N. (2011). Frustration rising: IBM 2011 commuter pain survey. *IBM, NY, USA*.

Harrison, P., Gotz, G., Todes, A., Wray, C., Ahmed, P., & Badenhorst, W. (2010). *Changing space, changing city: Johannesburg after apartheid*. Wits University Press.

Hazlett, C. & Xu, Y. (2018). Trajectory Balancing: A General Reweighting Approach to Causal Inference With Time-Series Cross-Sectional Data. *SSRN Electronic Journal*.

Holzer, H. J., Quigley, J. M., & Raphael, S. (2003). Public transit and the spatial distribution of minority employment: Evidence from a natural experiment. *Journal of Policy Analysis and Management*, *22*(3), 415–441.

Jang, S., An, Y., Yi, C., & Lee, S. (2017). Assessing the spatial equity of Seoul's public transportation using the Gini coefficient based on its accessibility. *International Journal of Urban Sciences*, *21*(1), 91–107.

Johannesburg Strategic ITP Framework. (2013). *Strategic integrated transport plan framework for the City of Johannesburg*.

Johnson, V., Currie, G., & Stanley, J. (2011). Exploring transport to arts and cultural activities as a facilitator of social inclusion. *Transport Policy*, *18*(1), 68–75.

Kain, J. F. (1968). Housing Segregation, Negro Employment, and Metropolitan Decentralization. *The Quarterly Journal of Economics*, *82*(2), 175.

Kaul, A., Klößner, S., Pfeifer, G., & Schieler, M. (2015). Synthetic control methods: Never use all pre-intervention outcomes together with covariates.

Kennan, J. & Walker, J. R. (2011). The effect of expected income on individual migration decisions. *Econometrica*, *79*(1), 211–251.

Kenyon, S. (2006). The 'accessibility diary': Discussing a new methodological approach to understand the impact of internet use upon personal travel and activity participation. *Journal of Transport Geography*, *14*(2), 123–134.

Kerr, A. (2017). Tax(i)ing the Poor? Commuting Costs in South African Cities. *South African Journal of Economics*, *85*(3), 321–340.

Kishor, S. & Neitzel, K. (1996). The status of women: Indicators for twenty-five countries.

Kumar, A., Zimmerman, S., & Agarwal, O. (2012). The soft side of brt: Lessons from five developing cities. *The International Bank for Reconstruction and Development/The World Bank Group*.

Laird, J. J. & Venables, A. J. (2017). Transport investment and economic performance: A framework for project appraisal. *Transport Policy*, *56*, 1–11.

Lall, S. V., Henderson, J. V., & Venables, A. J. (2017). *Africa's cities: Opening doors to the world*. The World Bank.

Lehohla, P. (2003). New geographical hierarchy structure. Retrieved May 11, 2019, from https://www.iol.co.za/business-report/opinion/new-geographical-hierarchy-structure-773527

Levitas, R., Pantazis, C., Fahmy, E., Gordon, D., Lloyd, E., & Patsios, D. (2007). The Multidimensional Analysis of Social Exclusion, 247.

L'Hour, J. & Abadie, A. (2018). A Penalized Synthetic Control Estimator for Disaggregated Data, 49.

Litman, T. (2014). Evaluating Transportation Equity, 53.

Lucas, K. (2011). Making the connections between transport disadvantage and the social exclusion of low income populations in the Tshwane region of South Africa. *Journal of Transport Geography*, *19*(6), 1320–1334.

Lucas, K. (2012). Transport and social exclusion: Where are we now? *Transport Policy*, *20*, 105–113.

Machard, D. C. (2014). *A case study of the socio-economic impacts of the school commute within the inner City of Johannesburg* (Doctoral dissertation, University of Johannesburg).

Majid, H., Malik, A., & Vyborny, K. (2018). Infrastructure investments and public transport use: Evidence from Lahore, Pakistan, 56.

McClelland, R. (2017). The Synthetic Control Method as a Tool to Understand State Policy, 46.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of Public Economics*, *3*(4), 303–328.

McLafferty, S. (1997). Gender, race, and the determinants of commuting: New York in 1990. *Urban Geography*, *18*(3), 192–212.

Melo, P. C., Graham, D. J., & Brage-Ardao, R. (2013). The productivity of transport infrastructure investment: A meta-analysis of empirical evidence. *Regional Science and Urban Economics*, *43*(5), 695–706.

Mlatsheni, C. & Leibbrandt, M. (2014). Unemployment in South Africa. In H. Bhorat, A. Hirsch, R. Kanbur, & M. Ncube (Eds.), *The Oxford Companion to the Economics of South Africa* (pp. 236–243). Oxford University Press.

Mlatsheni, C. & Ranchhod, V. (2017). Youth labour market dynamics in South Africa: Evidence from NIDS 1-2-3, 28.

Monte, F., Redding, S. J., & Rossi-Hansberg, E. (2018). Commuting, Migration, and Local Employment Elasticities. *American Economic Review*, *108*(12), 3855–3890.

Municipal Demarcation Board. (2016). Ward delimitation. Retrieved May 11, 2019, from http://www.demarcation.org.za/site/ward-delimitation/

Municipal Demarcation Board. (2019). Municipal ward boundaries. Retrieved May 11, 2019, from http://www.sasdi.net/metaview.aspx?uuid=2b9efabcfb69dc4a21f71626a371c4be

Ngobeni, J. M. (2014). Migrant Characteristics and Migration Patterns to Gauteng (2001-2011), 40.

Observatory, G. C.-R. (2019). Gauteng wards 2009.

OECD. (2011). Work and life balance. OECD.

OECD. (2019a). Unemployment rate (indicator). Retrieved May 10, 2019, from https://data.oecd.org/unemp/unemployment-rate.htm

OECD. (2019b). Youth unemployment rate (indicator). Retrieved May 10, 2019, from https://data.oecd.org/unemp/youth-unemployment-rate.htm

Onatu, G. O. (2011). A New Approach to Sustainable Transport System: A Critical Appraisal of the Rapid Bus Transit System (BRT) in South Africa., 9.

Picarelli, N., Jaupart, P., Chen, Y., & Moreno-Monroy, A. (2017). Unequal Commutes. Job Accessibility & Employment in Accra, 52.

Powell, D. (2017). Synthetic Control Estimation Beyond Case Studies: Does the Minimum Wage Reduce Employment? 46.

Redding, S. J. & Turner, M. A. (2015). Transportation costs and the spatial organization of economic activity. In *Handbook of Regional and Urban Economics* (Vol. 5, pp. 1339–1398). Elsevier.

Risimati, B. & Gumbo, T. (2018). Exploring the applicability of location-based services to delineate the state public transport routes integratedness within the City of Johannesburg. *Infrastructures*, *3*(3), 28.

Rogerson, C. & Rogerson, J. (1997). The Changing Post-apartheid City: Emergent Black-owned Small Enterprises in Johannesburg. *Urban Studies*, *34*(1), 85–103.

Rotger, G. A. P. & Nielsen, T. A. S. (2015). Effects of job accessibility improved by public transport system: Natural experimental evidence from the Copenhagen metro. *European Journal of Transport and Infrastructure Research*, *15*(4), 419–441.

Samii, C. (2019a). Inverse covariance weighting versus factor analysis. Retrieved May 11, 2019, from http://cyrussamii.com/?p=2177

Samii, C. (2019b). Makeindex project. Retrieved May 11, 2019, from https://github.com/cdsamii/make_index

Scorcia, H. & Munoz-Raskin, R. (2019). Why South African cities are different? Comparing Johannesburg's Rea Vaya bus rapid transit system with its Latin American siblings. *Case Studies on Transport Policy*.

Selassie, A. A. (2011). South Africa's unemployment puzzle. *IMFDirect Blog*.

Selod, H. & Zenou, Y. (2001). Location and Education in South African Cities under and after Apartheid. *Journal of Urban Economics*, *49*(1), 168–198.

Statistics South Africa. (2004). *Census 2001: Primary tables South Africa: Census '96 and 2001 compared* (Stat SA Library Cataloguing-in-Publication (CIP) Data No. 03-02-04 (2001)). Statistics South Africa. Pretoria, South Africa.

Statistics South Africa. (2011). Average household income. Retrieved May 10, 2019, from http://www.statssa.gov.za/?page_id=4286&id=11317

Statistics South Africa. (2014). National household travel survey.

Statistics South Africa. (2019a). Census 1996, 2001, 2011. Retrieved May 11, 2019, from http://superweb.statssa.gov.za/webapi/jsf/dataCatalogueExplorer.xhtml

Statistics South Africa. (2019b). Main place areas. Retrieved May 11, 2019, from http://www.sasdi.net/metaview.aspx?uuid=8942a0d3989367e1421ebb644d3d03f1

Tsivanidis, N. (2018). The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogota's TransMilenio, 62.

Turner, T. & Niemeier, D. (1997). Travel to work and household responsibility: New evidence. *Transportation*, *24*(4), 397–419.

van Ham, M., Mulder, C. H., & Hooimeijer, P. (2001). Spatial Flexibility in Job Mobility: Macrolevel Opportunities and Microlevel Restrictions. *Environment and Planning A: Economy and Space*, *33*(5), 921–940.

Vaz, E. & Venter, C. (2012). The Effectiveness of Bus Rapid Transit as Part of a Poverty-Reduction Strategy: Some Early Impacts in Johannesburg, 14.

Venter, C. & Hayes, G. (2017). South Africa needs to revamp its new public transport system. The Conversation.

Venter, C., Jennings, G., Hidalgo, D., & Valderrama Pineda, A. F. (2018). The equity impacts of bus rapid transit: A review of the evidence and implications for sustainable transport. *International Journal of Sustainable Transportation*, *12*(2), 140–152.

Vuri, D. (2010). The Effect of Availability of School and Distance to School on Children's Time Allocation in Ghana: The Effect of Availability of School and Distance to School. *LABOUR*, *24*, 46–75.

Wachs, M. & Kumagai, T. (1973). Physical accessibility as a social indicator. *Socio-Economic Planning Sciences*, *7*(5), 437–456.

Woolf, S. & Joubert, J. (2013). A people-centred view on paratransit in South Africa. *Cities*, *35*, 284–293.

Xu, Y. (2017). Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models. *Political Analysis*, *25*(1), 57–76.

Zax, J. S. & Kain, J. F. (1996). Moving to the suburbs: Do relocating companies leave their black employees behind? *Journal of Labor Economics*, *14*(3), 472–504.

Zimran, A. (2018). *Transportation and Health in a Developing Country: The United States* (tech. rep. No. w24943). National Bureau of Economic Research. Cambridge, MA.

9 Appendix

Predictor	Treated	Synthetic
Unemployment	.4906555	.4944207
Female-headed HH	.4250081	.359324
Fraction of white	.0016811	.0001836

Table 8: Predictor balance (raw)

Donor Units	Weight
1	.029
2	.971
4	0
5	0

Table 9: Donor weights (raw)

RMSPE	.032436	.032544	.032544	.032991	.035759	.035873	.039981	.043055
Unemployment	x	x	x	x		x	x	x
Female headed HH	x			x				x
Fraction of white	x					x	x	x
High school ed.		x	x		x		x	
HH facility index		x		x	x	x	x	x

Table 10: Pre-treatment RMSPEs by specification (unemployment raw)

Ward Number	Distance	Ward Number	Distance	Ward Number	Distance
11	1.7	28	1.3	44	2.3
12	.2	30	.9	45	1.1
13	.8	31	.5	46	.6
14	.8	33	.2	47	.3
15	.6	34	1	48	1.6
16	1.3	35	.3	49	>3
19	.7	36	.4	50	>3
20	.3	37	.8	51	1.4
21	.2	38	.3	52	1.4
22	1.2	39	.6	53	>3
24	2.3	40	.8	128	>3
25	.4	41	.4	129	>3
26	2.3	42	.2	130	.8
27	1	43	1.1		

Table 11: Treatment intensities by ward, kilometer distance from nearest BRT stop

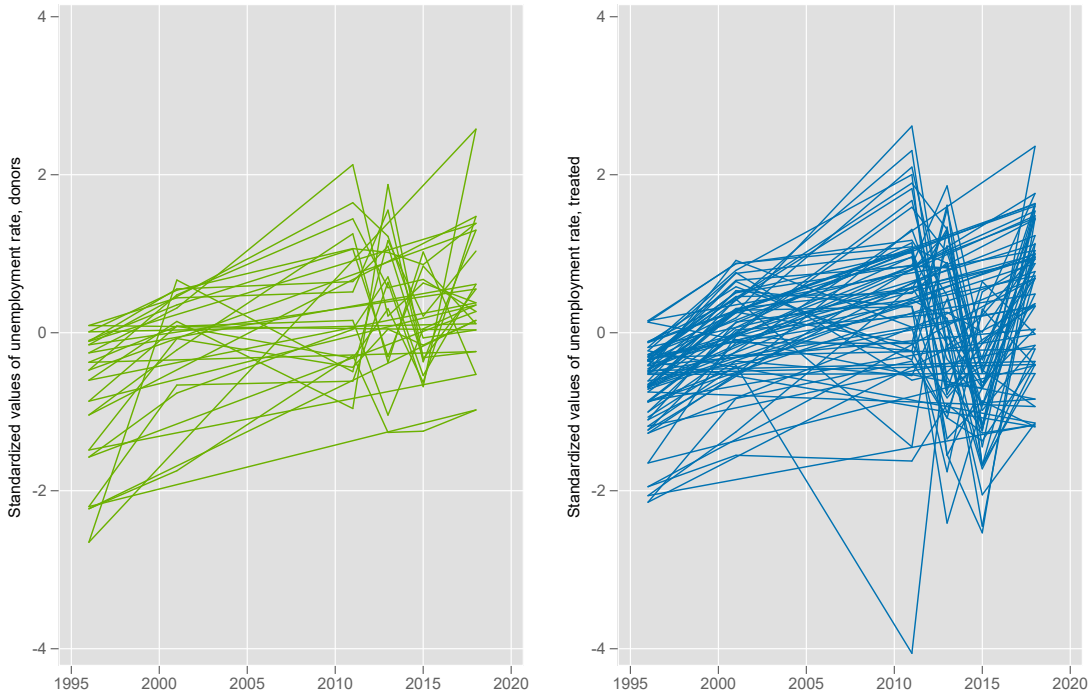


Figure 19: Unemployment rate - all wards

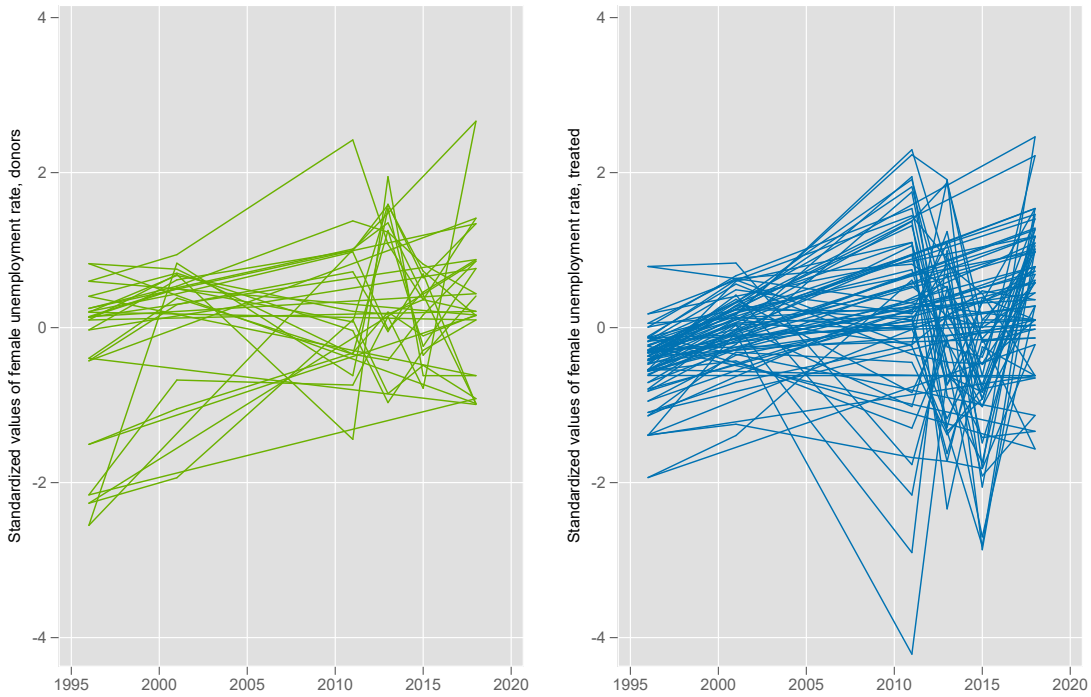


Figure 20: Female unemployment rate - all wards

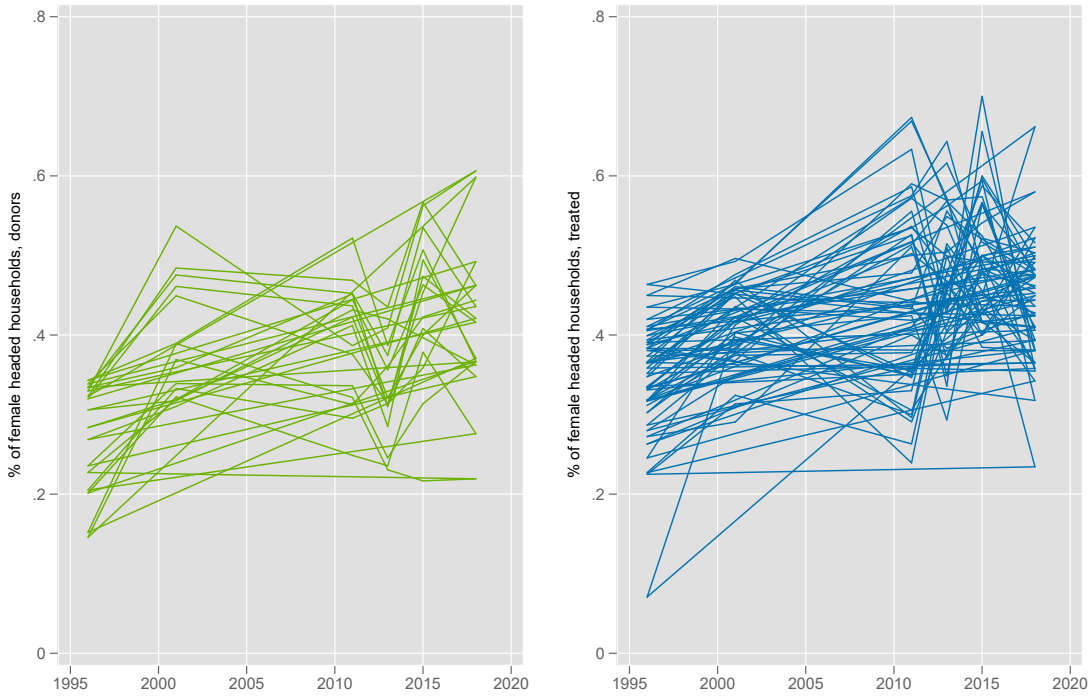


Figure 21: Female headed households - all wards

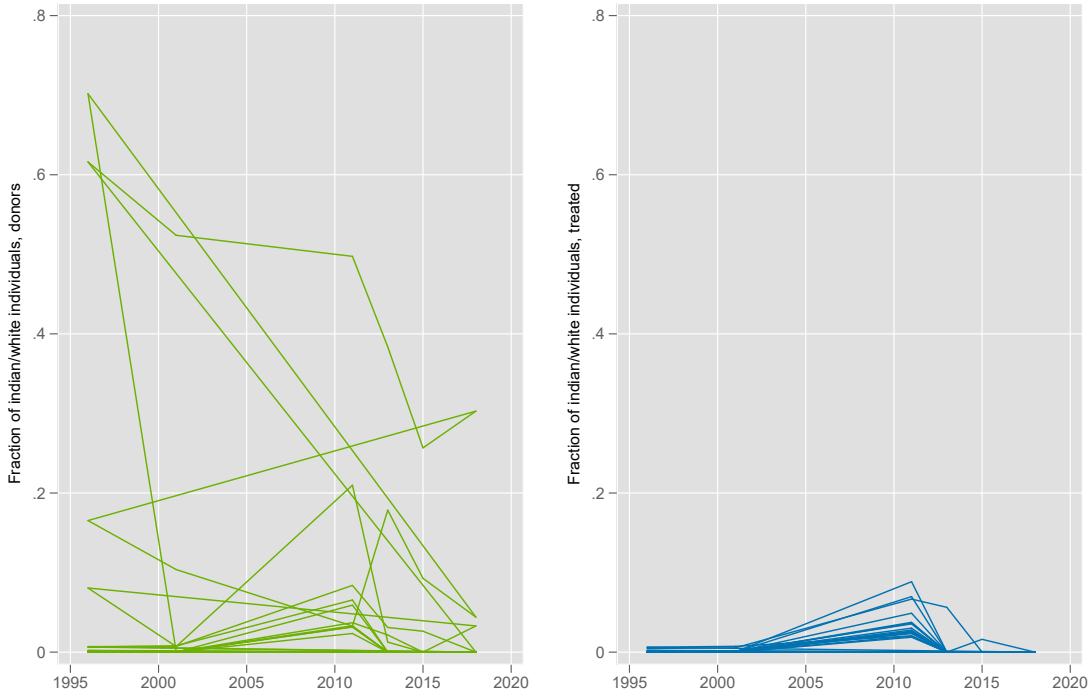


Figure 22: Fraction of Indian/Asian/white households - all wards

Unit	1996	2001	2011	2013	2015	2018
Group 1	-.0140346	.0145799	-.0898968	-.0823635	-.1585275	-.0588523
Group 2	-.007028	.0100822	.0435686	-.0730755	-.2103962	-.038408
Group 3	-.0162513	.0162838	-.0164969	-.1142458	-.0947019	-.0356677
Donor 1	.0315133	-.0320829	.1344968	-.150196	.1378019	-.1235286
Donor 2	.0227375	-.022723	.1096762	.0611886	.1728962	.0651759
Donor 3	-.0032672	-.0019755	-.0246209	.0450787	-.1128422	-.0788969
Donor 4	.0144119	-.0134688	.1577484	.0844446	-.0035876	-.113014
Donor 5	.0291726	-.0282402	-.02577	.1622317	.0090814	.085963
Donor 6	-.0030806	.0039184	.1188972	.0608007	.1530266	-.0390253
Donor 7	-.0024226	.0063386	-.0713622	.1246055	.1420209	.136521
Donor 77	.0288201	-.0284393	.279744	.0542085	.1219478	-.0059596
Donor 78	.0156294	-.0146139	-.1538222	.2873635	.0644682	-.0269758
Donor 79	.0365913	-.0355821	.0331156	-.0804598	.0804798	.0310779
Donor 80	-.0031374	-.0006814	.0347264	.1032652	-.01607	.1913177
Donor 95	.0016533	-.0011715	.259026	.0243795	.2134141	-.040517
Donor 113	-.0372146	.0352696	.0115207	.0640395	.0678335	.0319838
Donor 120	-.0356071	-.0249573	.1703745	-.0457746	.002935	.0241917
Donor 121	-.059216	.0626744	.0584499	.1951371	.0467244	.4317537

Table 12: Treatment effects by year

Unit	1996	2001	2011	2013	2015	2018
Group 1	-0.01033	0.009855	-0.01272	-0.14716	-0.133	-0.00686
Group 2	-0.00409	0.003093	0.062726	-0.1228	-0.26357	-0.02159
Group 3	0.000935	0.000675	-0.06351	-0.10336	-0.11296	0.062718
Donor 1	0.011792	-0.00931	0.170957	-0.16191	0.01668	-0.02668
Donor 2	0.009811	-0.00982	-0.09229	0.117745	0.131897	0.170635
Donor 3	0.002398	-0.00255	0.203754	-0.00313	0.054831	-0.0823
Donor 4	-9.3E-05	-0.00108	0.186157	0.056696	0.028868	-0.26568
Donor 5	0.019165	-0.0191	-0.11121	0.101765	-0.0576	0.088571
Donor 6	-0.02094	0.018513	0.125368	0.161946	0.154845	-0.054
Donor 7	0.003979	-0.00223	-0.14958	0.066228	0.227723	0.102918
Donor 77	-0.02755	0.01621	0.4916	0.085735	0.064019	-0.0224
Donor 78	-0.0198	0.023231	-0.09569	0.274335	-0.04275	-0.09881
Donor 79	-0.00119	0.005762	0.216197	-0.12974	-0.00068	-0.01923
Donor 80	-0.02139	0.024588	0.211606	0.058089	0.096386	0.191574
Donor 95	-0.03809	0.040464	0.295612	0.082689	0.198406	-0.14137
Donor 113	-0.02122	0.021062	0.070793	-0.04114	0.085175	0.092969
Donor 120	-0.04611	-0.0758	-0.07136	-0.09751	0.198519	0.030154
Donor 121	-0.09052	0.090731	0.072337	0.222089	0.090591	0.422665

Table 13: Treatment and placebo effects by year, female unemployment

Treatment Wards					
Ward	**ATE**	**P-Values**	**Ward**	**ATE**	**P-Values**
11	-.0810971	.7333333	37	-.0728524	.5333334
12	-.1692558*	0	38	-.0823597	.2666667
13	-.1234222	.2666667	39	-.1376801*	.0666667
14	-.096537*	.0666667	40	.0158942*	.0666667 .
15	-.1240868	.1333333	41	-.1302558	.2
16	-.1534297*	.0666667	42	-.1353599*	.0666667
19	-.0534528*	.0666667	43	-.1844327*	.0666667
20	-.0586825	.6666667	44	.0108838	.3333333
21	-.0555645	.3333333	45	.0043434	.3333333
22	-.0579527	.4	46	-.0654052	.3333333
24	-.0196528	.8666667	47	-.0644041	.6666667
25	-.1222424*	.0666667	48	-.0144952	.8666667
26	-.0747184	.4	49	-.1292472	.2
27	-.1034353*	.0666667	50	-.0980771*	.0666667
28	-.0056087*	.0666667	51	-.0993955	.3333333
30	-.0226924	.8666667	52	-.0309808	.6666667
31	.0273205	.7333333	53	-.1622028*	.0666667
33	-.1690177*	.0666667	128	.0728059	.2
34	.0284303	.2	129	-.0188958	.6666667
35	-.0786648*	.0666667	130	-.07969363	.4
36	-.1304335	.2666667			

Table 14: Average treatment effects and significance, unemployment

Treatment Wards					
Ward	**ATE**	**P-Values**	**Ward**	**ATE**	**P-Values**
11	-.0528079	.9333333	37	-.1007364	.2
12	-.1858071*	0	38	-.0226738	.9333333
13	-.2202313	.1333333	39	-.1360945*	0
14	-.0574898	.4666667	40	.0686955	.4666667
15	-.2295652*	0	41	-.0519381	.2666667
16	-.1854201	.1333333	42	-.0757654	.1333333
19	-.0300094*	0	43	-.1766314*	0
20	-.0512935	.9333333	44	.0495887	.4666667
21	-.0929736	.2666667	45	-.0685837	.1333333
22	-.0987982	.2666667	46	-.0753186	.7333333
24	.0640905	.2	47	-.1174673	.2
25	-.0322872	.8666667	48	-.022131	1
26	-.0499514	.4666667	49	-.2069739*	0
27	-.1116406	.1333333	50	-.1048298	.1333333
28	.0237213	.4666667	51	-.0438622	.4
30	-.0716546	.8666667	52	.0444692	.8666667
31	-.0244814	.6	53	-.0900316	.3333333
33	-.1550704*	0	128	.0338681	.4666667
34	.0916234*	0	129	-.0754633	.8666667
35	-.0711392*	.0666667	130	-.1632608	.2
36	-.0182863	.9333333			

Table 15: Average treatment effects and significance, female unemployment

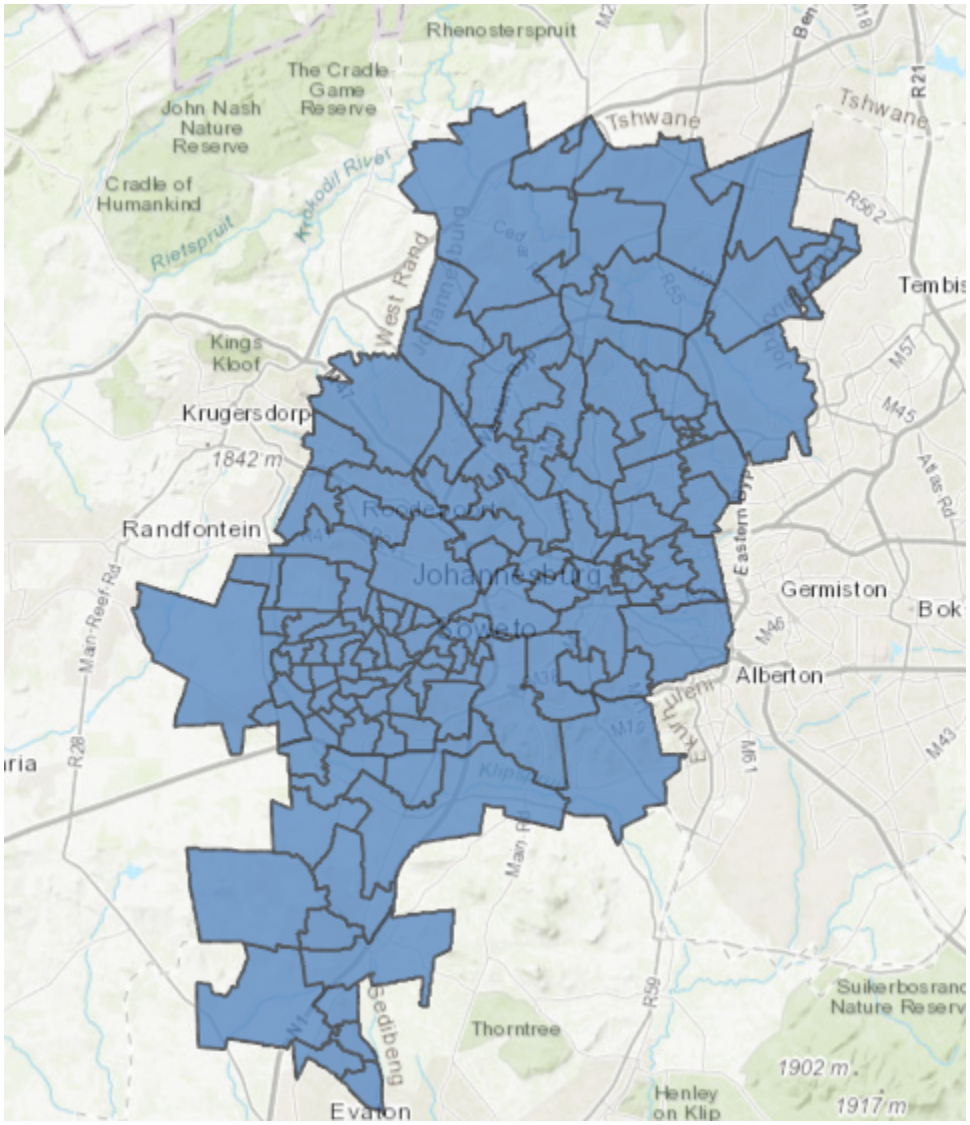


Figure 23: City of Johannesburg electoral wards (2011)
Source: Authors' illustration of Municipal Demarcation Board maps

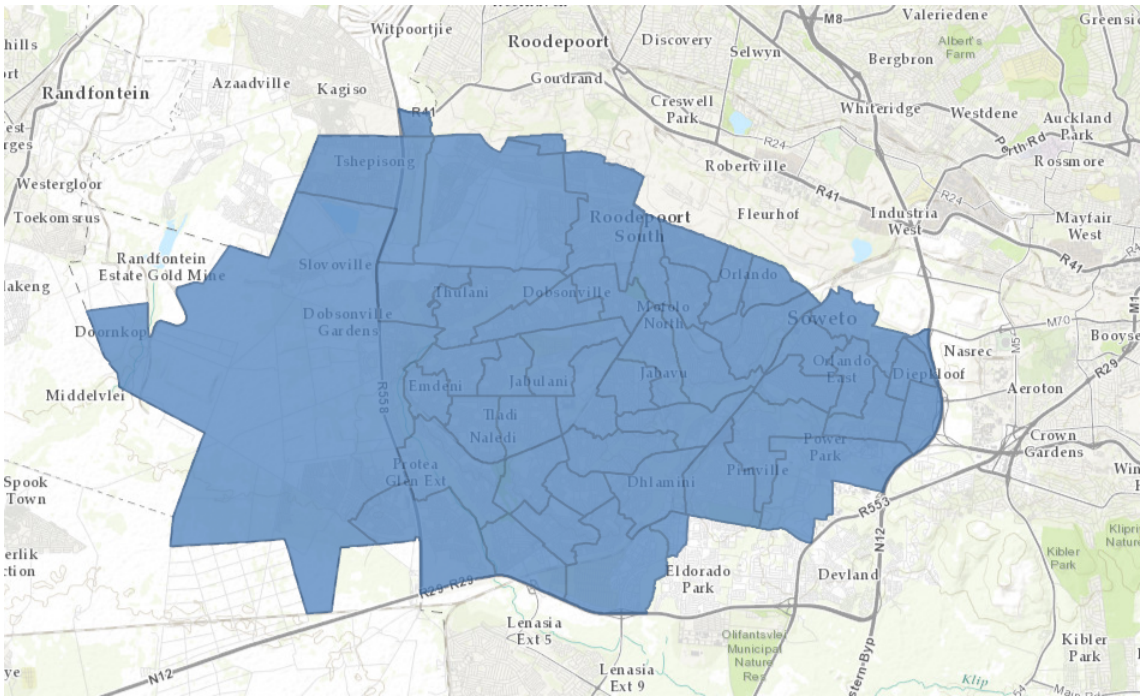


Figure 24: Electoral wards in Soweto (2011)
Source: Authors' illustration of Municipal Demarcation Board maps

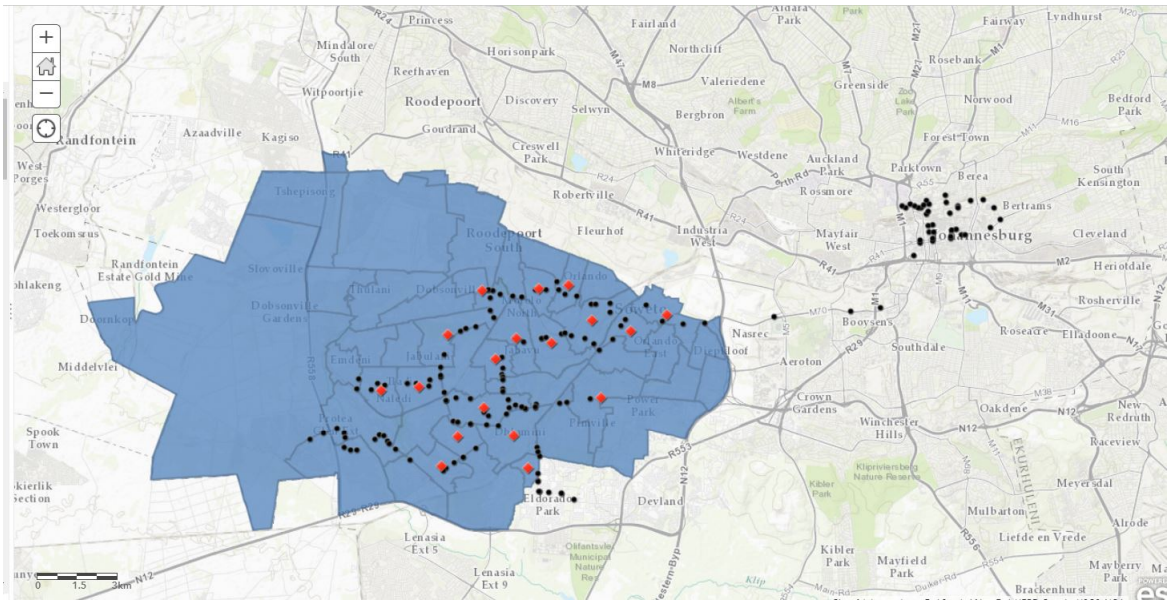


Figure 25: Rea Vaya phase 1a route and stops with 0.7 km ward centroids (Soweto)
Source: Authors' illustration of Municipal Demarcation Board maps

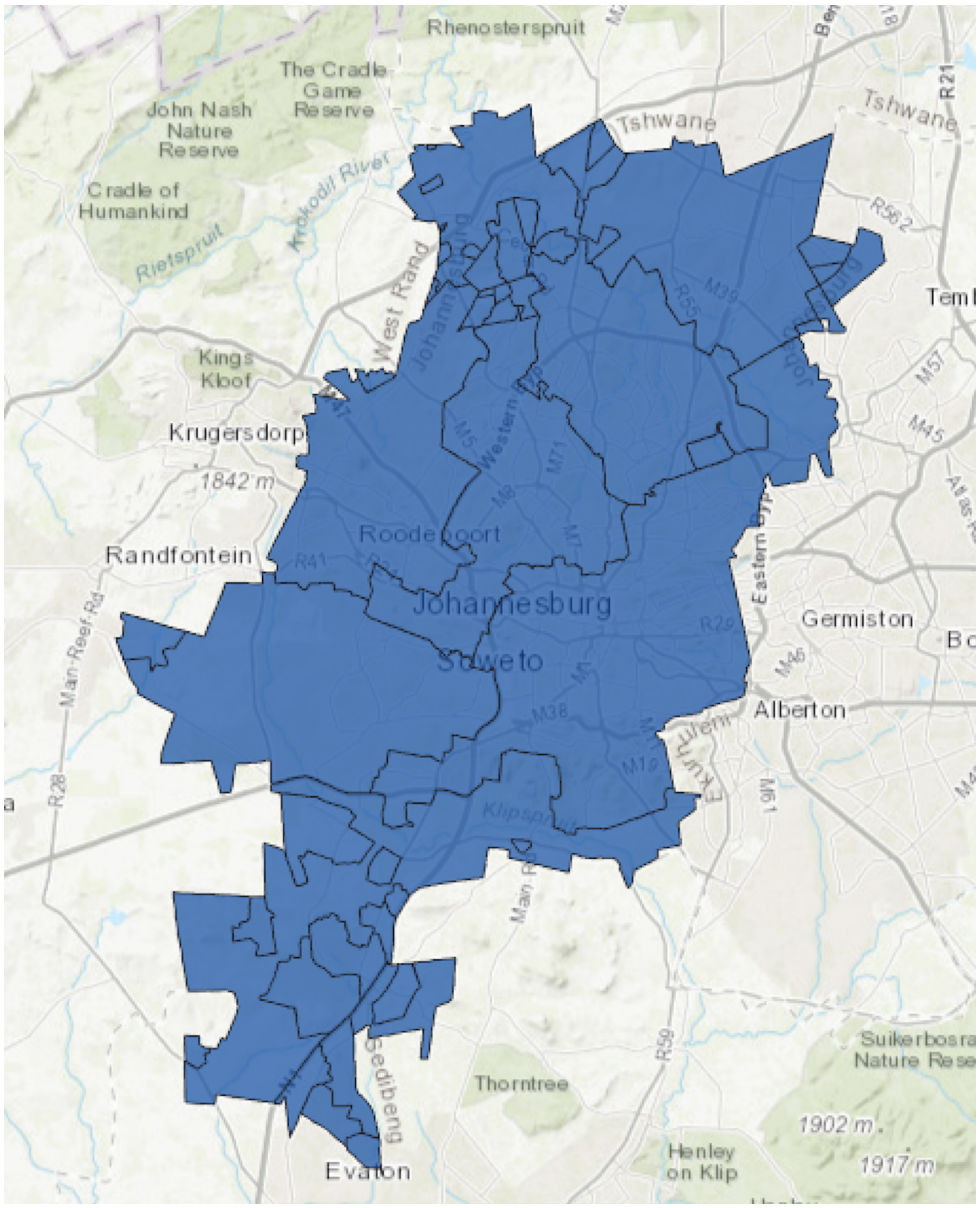


Figure 26: Main places in COJ

Source: Authors' illustration of Municipal Demarcation Board maps

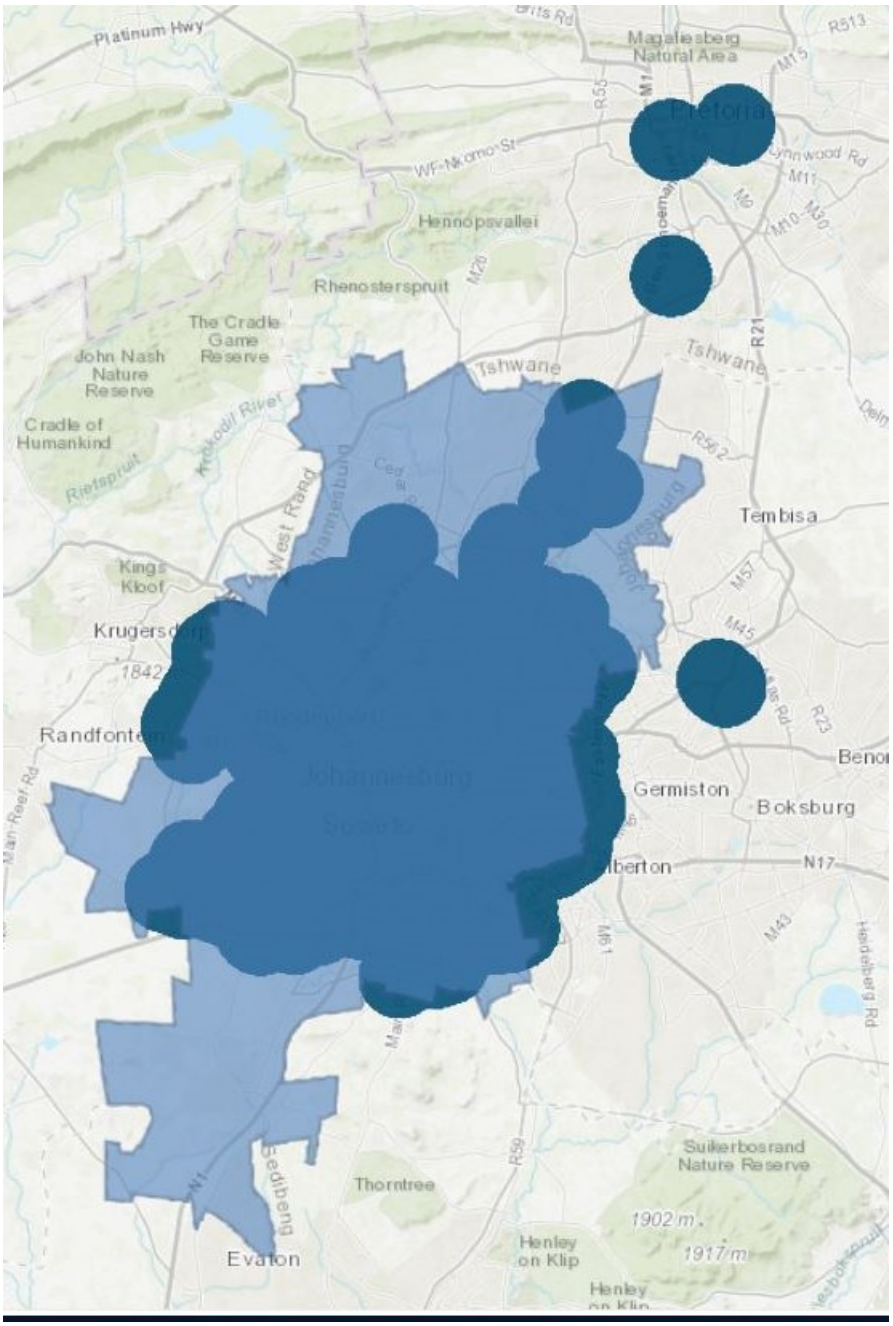


Figure 27: COJ with 3 km exclusion buffer of treatment
Source: Authors' illustration of Municipal Demarcation Board maps

FIGURE 5: THE MINIBUS-TAXI ROUTE NETWORK IN THE CITY OF JOHANNESBURG (ALL OPERATORS)

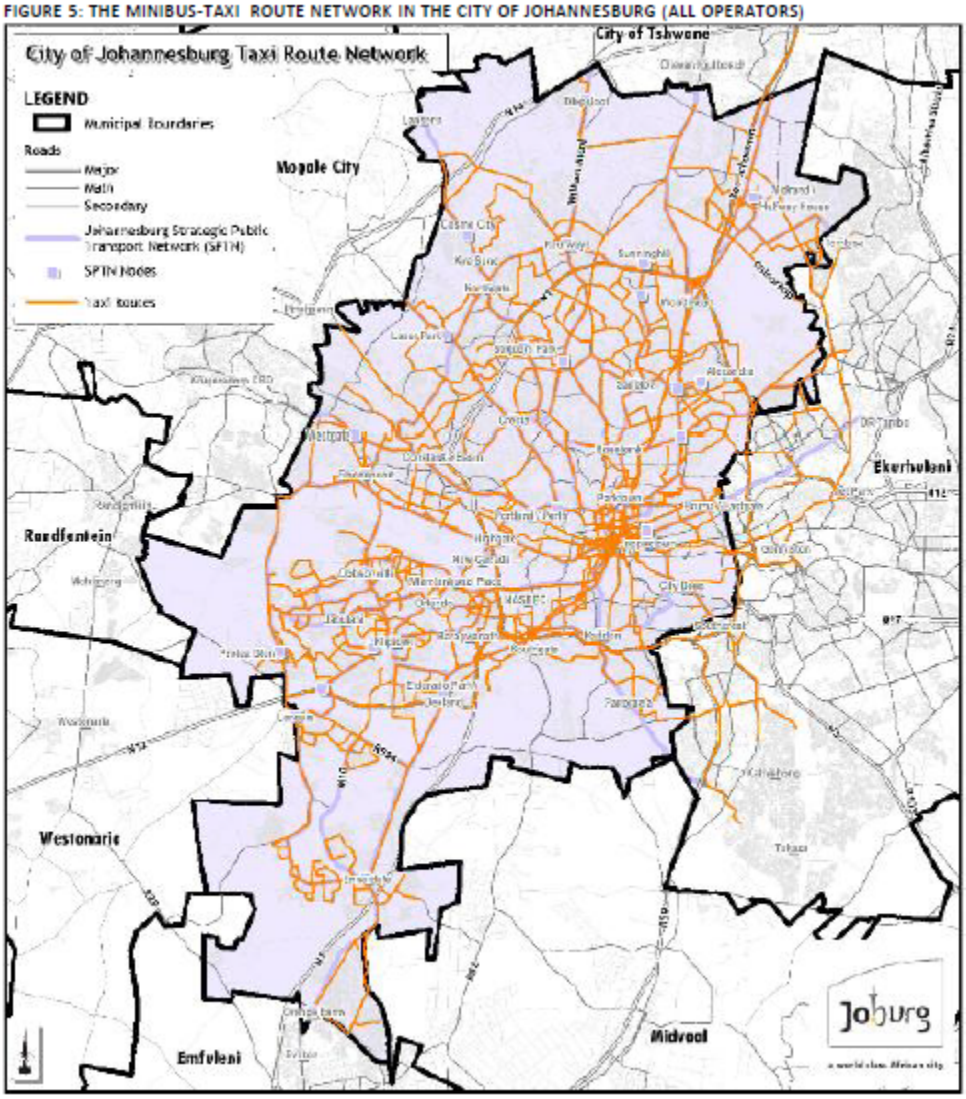


Figure 28: Taxi network in COJ
Source: City of Johannesburg