CODERA ANALYTICS

Is Your Metro Hung-Up? Labour Market Effects of Stable Tenure in a South African Metro

Jacques Quass de Vos[1]

POLICY PAPER 26-8
August 2026

Abstract

How much of a city's economic performance can be attributed to the form and performance of its leadership, and how much is forgone amid the instability and political competition that increasingly characterise South African local government? Using municipal-level employment data and a synthetic difference-in-differences approach, this paper estimates the labour market effects of an incumbent mayoral tenure in a South African metropolitan municipality. While inference is constrained by the single treated unit, the estimates suggest that a stable, single-party mayoral tenure in Cape Town may have raised the labour force absorption rate by between 2.3 and 5.5 percentage points, depending on specification.

JEL classification: R23, H7, D72, C23

Keywords: synthetic difference-in-differences, political stability, coalition, local government, employment, migration, South Africa

Acknowledgements: I would like to thank my colleagues at Codera Analytics, Daan Steenkamp and Aidan Horn, for their helpful comments and prior analysis on related topics. I am further indebted to Chen Lin and Taira Hasegawa for their valuable feedback. Remaining mistakes are my own. The views in this paper are solely the responsibility of the author and do not reflect the views of any affiliated institutions. The findings should not be read as an endorsement of any political party. While the author was previously employed by the Western Cape Government, this was in a non-partisan administrative capacity.

[1] Codera Analytics. The University of Osaka, Graduate School of Economics.

1 Introduction

Following South Africa's first democratic elections in 1994, the African National Congress (ANC) has enjoyed a near-monopoly over government at the national, provincial, and local levels. ANC majority has most prominently been challenged by the Democratic Alliance (DA) whose outright control centres on the Western Cape. More recently, sustained dissatisfaction with the ANC has eroded its dominance and lifted support for opposition parties. Yet the DA's support base remains heavily concentrated among minority voters (Ferree, 2006; Paret and Runciman, 2023), and no opposition party has proved able to command outright majorities at scale. The result is a proliferation of municipalities governed by coalition (or supply-and-confidence) arrangements.

This paper empirically investigates what the ensuing instability costs. We focus on Cape Town, which following the November 2021 election entered a sustained period of stable, single-party majority governance under Mayor Geordin Hill-Lewis (henceforth GHL), in contrast both to its own recent past and to the turmoil affecting other metros over the same period. Treating the onset of this administration as a discrete intervention, we use a synthetic difference-in-differences design to estimate its effect on the metro's labour force absorption rate, constructing a synthetic Cape Town from the other seven metros. The estimated effect is positive across every aggregate specification, and is robust to a series of placebo tests. Industry-level estimates on full-time employment are also mostly positive. What is less clear is the magnitude of the effect. Our preferred estimate of 5.5 percentage points is taken as an upper bound. A stressed specification that moves treatment one year later, deliberately biasing the estimate downward, yields a lower bound of 2.3 percentage points.

As of the 2016 local elections, four of South Africa's eight metropolitan municipalities were governed by coalition; by the November 2021 elections this had risen to five, with the number of hung councils among local municipalities more than doubling (IEC, 2021). The May 2024 general election then produced, at national level, a ten-party Government of National Unity (GNU). Coalition governance need not be dysfunctional, but in the South African case—absent a robust framework to structure such rule—it has frequently produced budget impasses, antagonistic council politics, and the stalling of new policy, with local government treated as a political battleground rather than a vehicle for service delivery (Tladi, 2024).

Coalition fragility is not the only source of instability. Majority-governed metros have proved similarly prone to turmoil generated from within the ruling party through tightly contested motions of no confidence, factional removals, and the mid-term loss of mayors. Cape Town's own recent past is a good example: Patricia de Lille's protracted removal in 2018, and the departure of senior councillors that followed, mark a period of considerable turmoil under majority rule. The repeated intra-ANC removals in metros such as Nelson Mandela Bay and eThekwini make the same point. What most clearly distinguishes one metro from another, then, is less whether it is governed by coalition than whether its governance is stable.

A large body of literature examines individual leadership and its effects on economic and business outcomes. Jones and Olken (2005), using an original dataset, exploit the natural and accidental deaths of heads of state as a source of exogenous variation to test whether national leadership—and its quality—affects economic growth. Besley et al. (2011) extend this dataset to show that the educational attainment of rulers also matters: when a highly educated leader is succeeded by a relatively unlearned one, the economic losses are stark. Bertrand and Schoar (2003) ask whether individual top executives matter for corporate policy and firm performance, finding that manager fixed effects account for significant variation across a wide range of decisions and financial results. Yao and Zhang (2015), in a setting

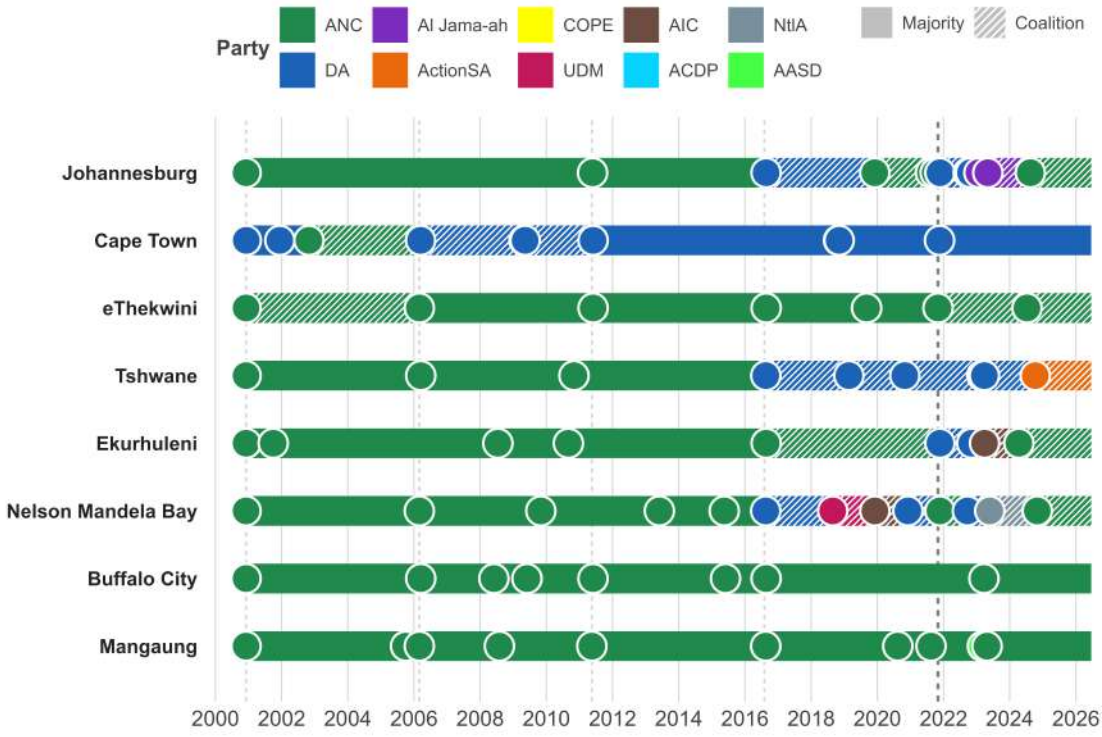

Figure 1: Timeline of mayoral handovers in South African metros, 2000–2026

closest to our metro case, estimate the personal effects of mayors and party secretaries on the growth of Chinese cities. By leveraging the homogeneous institutional characteristics provided by the ruling Chinese Communist Party on a national scale, they are able to isolate individual leader effects, arguing that better leadership may promote better growth.

Beyond individual effects, the structure and composition of government also shape policy-making and a government's capacity to rule effectively. Roubini and Sachs (1989) find that shorter tenures and the presence of multiple parties in a ruling coalition are correlated with larger fiscal deficits. Alesina and Drazen (1991) provide a game-theoretic basis for studying coalition governments, arguing that political stabilisation resembles a war of attrition.[1] In a regression discontinuity design, Ferreira and Gyourko (2009) compare U.S. cities where a Democratic mayoral candidate narrowly won to those where one narrowly lost, finding that partisanship has no significant effect on outcomes such as local government size, crime rates, or the composition of budgets. Taken together, this work suggests that local economic outcomes cannot be reduced to deterministic structural forces alone: the individuals who run governments, and the stability of the arrangements through which they do so, plausibly shape what those governments achieve.

Despite this, the economic effects of governance stability at the subnational level remain largely unexamined in the South African context, and the existing leadership literature elsewhere has focused on partisanship or individual identity rather than tenure stability per se. Simultaneously, as the 2026 local government elections draw near, there is an increased urgency amongst researchers and political commentators pushing for clearer legislation related to coalition governance in municipalities.[2] This paper hopes to contribute to the ongoing

[1] Factions delay concession to shift the costs of adjustment onto their rivals, so that fragmented or unstable governments systematically under-supply timely policy.

[2] See, e.g., de Vos (2021) for a chapter on the constitutional and legal architecture governing coalition formation in South Africa, and why it leaves coalition and minority governments particularly vulnerable to politicking

debate by providing empirical evidence as to the stability mechanism enabling local labour markets to outperform their less stable counterparts.

The remainder of the paper is structured as follows. Section 2 sets out the data and methodology, together with our identification strategy and the threats to it; Section 3 presents the results; Section 4 assesses their robustness through a series of placebo tests; Section 5 decomposes the results by industry and then by industry groups; and Section 6 discusses the findings.

2 Data & Methodology

2.1 Data

Our primary outcome variable, the labour force absorption rate, is drawn from the Quarterly Labour Force Survey (QLFS), published every three months by Statistics South Africa (Stats SA). We use data from 2015–Q1 through 2026–Q1. We choose this measure for two reasons. First, in a country with unemployment above thirty percent and acute spatial inequality, employment is a central object of local political mandates: parties campaign on jobs, and metropolitan administrations are explicitly tasked with fostering the conditions for local economic activity. Absorption—the share of the working-age population in employment—is therefore a natural and policy-relevant measure of how well a metro's leadership performs against the goals on which it is elected and held to account.[3] Second, few alternative indicators remain that are publicly available, frequent, and spatially disaggregated to the metro level, making the QLFS particularly well suited to measuring metropolitan economic performance.

The QLFS master sample transitioned from a 2001- to a 2011-census frame in 2015–Q1. This change introduced a discontinuous and metro-heterogeneous shift in measured participation and unemployment (Kerr and Wittenberg, 2019). Starting in 2015–Q1 places the entire series on a single, consistent master frame, avoiding such a break in the data.[4] A more recent revision (the 2025–Q3 alignment of the QLFS to the latest ICLS resolutions) leaves the definitions of employment and unemployment, and hence the absorption rate, unchanged, affecting only the measurement of labour underutilisation and the sectoral classification of the informal economy. Our outcome is therefore comparable across the revision. We refer readers to von Fintel (2023)'s survey of metropolitan-level data availability and the relative strengths and weaknesses of the measures used to track South African cities' development.

As a robustness check, we later draw on the Spatial Tax Panel's full-time-equivalent (STP-FTE) employment series at the local-municipality level (Nell and Visagie, 2025b), to test whether the results hold across data sources and units of aggregation. The panel has monthly coverage from March 2013 to February 2025 and is based on SARS IRP5/PAYE administrative tax data. The data further disaggregate into industry classifications, which we later exploit to reveal industry-level dynamics. A drawback of the STP is the potential for head-office bias, whereby metros hosting firm headquarters may over-report employment when workers are in fact located at branches elsewhere (Nell and Visagie, 2025a). As of version 5.1 (we use version 6), a reallocation algorithm has been applied to mitigate this bias, though some cases

and instability; Mongale and Nkwatle (2025) review proposed amendments to the legislation governing local coalitions which appeared in a 21 May 2024 Government Gazette (CoGTA, 2024).

[3] Absorption rates are favoured over raw employment counts because, as a ratio rather than a level, the measure is not mechanically inflated by Cape Town's rapid population growth through semigration.

[4] The master frame remains consistent throughout our window; Stats SA had not transitioned the QLFS to a 2022-census frame as of 2026–Q1.

may go undetected. Encouragingly, the adjusted series aligns closely with the metro-level employment geography of the QLFS, supporting its use as an independent spatial check on our main results.

2.2 Estimation

Synthetic difference-in-differences (SDID), proposed by Arkhangelsky et al. (2021), estimates the causal effect of treatment by constructing a counterfactual for the treated unit from a weighted combination of control units and pre-treatment periods. The unit and time weights are chosen so that the reweighted donors track Cape Town's absorption rate over the pre-treatment period as closely as possible; the identifying assumption is that this parallel relation would have persisted into the post-treatment period absent intervention.[5] Measuring the difference between Cape Town's post-treatment outcome and that of its synthetic counterpart, after netting out their pre-treatment level offset, yields the average treatment effect on the treated (ATT).

SDID builds on the synthetic control tradition brought into the mainstream by Abadie and coauthors (Abadie and Gardeazabal, 2003; Abadie et al., 2010, 2015). It offers two further advantages in our setting. First, unlike the classical synthetic control method, it accommodates persistent level differences between metros by permitting the synthetic unit to differ from Cape Town by a constant. Second, it combines two sets of weights: unit weights, which align the donor pool with Cape Town's pre-treatment trend, and time weights, which down-weight pre-treatment periods least representative of the post-treatment window—guarding against bias from both unit- and time-specific confounders.

Below follows the SDID estimator equation:

$$\left(\hat{\tau}, \hat{\mu}, \hat{\alpha}, \hat{\beta}\right) = \underset{\tau,\mu,\alpha,\beta}{\arg\min} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \mu - \alpha_i - \beta_t - W_{it}\tau\right)^2 \hat{\omega}_i \hat{\lambda}_t \right\}. \tag{1}$$

In our application, Y_{it} is the outcome variable, i.e., the labour force absorption rate of metro i in quarter t. μ is a constant, α_i are metro fixed effects, which absorb persistent level differences such as Cape Town's structurally higher absorption base, and β_t are time fixed effects capturing common shocks across all metros in a given quarter (the COVID-19 disruption being the most salient). W_{it} indicates treatment, where $W_{it} = 1$ for Cape Town from 2022–Q3 onward and zero otherwise. The unit weights $\hat{\omega}_i$ and time weights $\hat{\lambda}_t$ are estimated prior to running Equation 1. $\hat{\tau}$ is the ATT over the post-treatment period, which is the variable we are interested in estimating.

2.3 Identification

Our identification rests on defining the mayoral term of Geordin Hill-Lewis as a "treatment" applied to the City of Cape Town.[6] The period prior to his first approved budget in July 2022 is taken to be one of sufficient instability to warrant an "untreated" status.[7] The remaining metros are likewise considered untreated for the full sample period: over this window, none

[5] The absence of treatment would imply that Cape Town suffered the same fate of coalition instability, mayoral handovers, and other severe obstacles to smooth governance as other metros in our sample.

[6] Incumbent as both Mayor of Cape Town and DA federal leader at the time of writing. GHL's accession to party leadership appears after our sample window, and thereby should not affect our estimate.

[7] We choose July 2022 as the date of treatment onset as this is the first budget enacted under the management of Mayor Hill-Lewis and, therefore, their first real opportunity to enact any meaningful decisions.

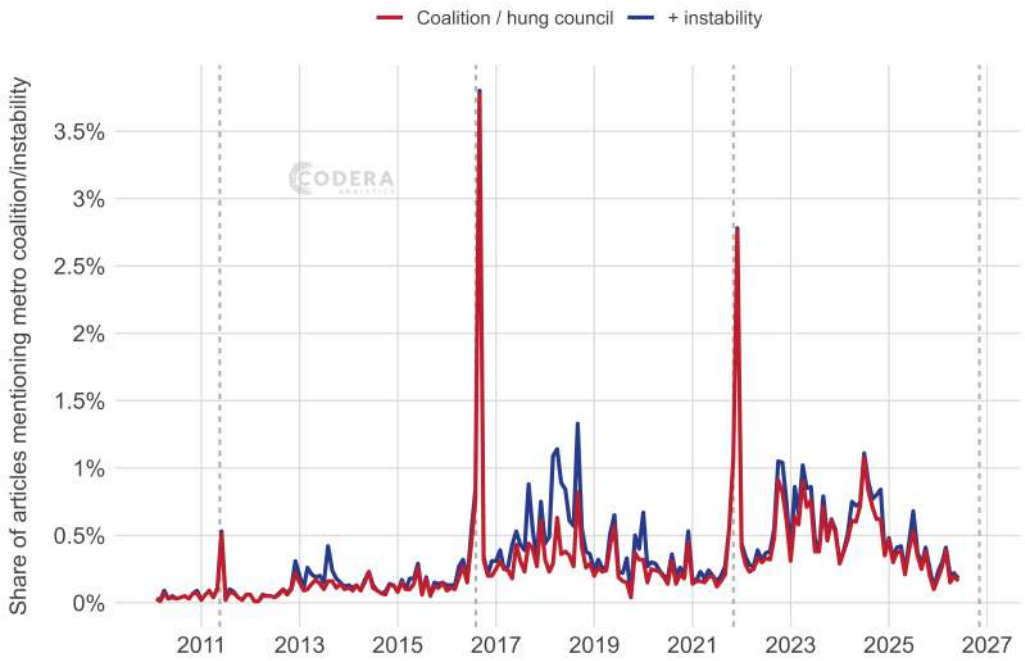

Figure 2: Monthly share of articles related to metropolitan coalition or instability

Notes: Data sourced from BusinessLive, BusinessTech, Daily Maverick, Mail & Guardian, and SABC online archives. A regular expression algorithm is used to identify the number of articles per month that contain terms related to metropolitan or local municipalities, coalitions, and council instability.

combined single-party majority rule with comparable governance stability, each being governed either through coalition arrangements or amid intra-party turmoil.

We reinforce our identification by reference to Figure 1, which illustrates the frequency of mayoral handovers since 2000 and the formation of coalitions after 2016. Figure 2 also lends support by showing how the share of South African online news articles related to municipal coalition talks and council instability has steadily increased since 2010.[8] Figure 3 reports the frequency with which South African metros are mentioned, per 10,000 words, within this municipal coalition and governance instability corpus.

At first glance, metros such as eThekwini, Buffalo City, and Mangaung might appear relatively settled, and comparable to the tenures of GHL or Patricia de Lille on the basis that they exhibit single-party majority and council consistency. On closer inspection, this is not the case. In eThekwini, Zandile Gumede faced corruption charges in May 2019 and announced a month later that she would step down as mayor, only to be formally removed by the ANC, alongside her co-accused councillors, in August of the same year (Tolsi, 2019; Grootes, 2019a,b).[9] Makhaye (2025) provides a brief summary of eThekwini's mayoral handovers and the related power struggles and controversies, evidence that eThekwini is sufficiently unstable across our study period.

Leadership transitions in Buffalo City have also often occurred during times of dire political fallout, such as former mayor Zukiswa Ncitha being accused of misappropriating funds intended for the late President Nelson Mandela's funeral in December of 2013 (Dayimani, 2022). Alfred Mtsi took over in 2015, only to be replaced by Xola Pakati a year later. In truth, Pakati's leadership was relatively stable considering he took over at a time when the metro

[8] Reassuringly, the index accurately tracks times of high coalition formation and reaches its highest peaks near local election periods.

[9] The initial charges implicated 62 city councillors.

suffered from a severe trust deficit. Yet, after serving for over seven years, he was recalled by the ANC over the city's alleged collapse, despite receiving unqualified audits and a 90% score from the National Treasury (Tandwa, 2023). Princess Faku, Pakati's successor and the incumbent, is serving a tenure that is growing increasingly unstable amid council dismissals and motions of no confidence.

Mangaung experienced a comparable intra-party removal: mayor Olly Mlamleli was unseated through a motion of no confidence in 2020 after surviving two earlier attempts, and the metro was placed under provincial administration in 2019 following National Treasury's intervention in its financial collapse (National Treasury, 2020). Lending further support to our identification are the 2021 and 2024 Governance Performance Indices (Good Governance Africa, 2021, 2024), where the donor metros generally score below Cape Town, which ranks best in both GPI years.[10] In summary, given the systemic volatility characteristic of metropolitan governance, the GHL administration represents a period of notable institutional stability that merits further empirical investigation.

2.4 Caveats and preliminary robustness discussion

There are several factors that may compromise the causal interpretation of our estimates, which we address in turn. First, as Bertrand and Schoar (2003) observe, there is no random allocation of top executives to firms; the same is plainly true of mayors elected to govern metros. GHL was not assigned to Cape Town as if by experiment; Cape Town is also the only DA-run metro in a DA-run province in our sample. Nevertheless, under the assumption that our estimator can produce an accurate counterfactual of Cape Town (i.e., parallel trends hold), we are able to recover a causal effect.

Second, it is possible that the treatment is mistimed. Since the DA has held an outright majority in Cape Town since 2011, the period preceding the November 2021 election may already have been sufficiently stable, meaning GHL's tenure cannot be regarded as a distinct intervention. We offer two responses: (1) The DA's ascent to a Cape Town majority was facilitated by the absorption of Patricia de Lille's Independent Democrats in 2010 and eventual complete merger in 2014; majority status was therefore not synonymous with settled single-party control from the outset. Given contention on core policy matters that would lead to friction between de Lille and DA leadership, we do not consider de Lille's leadership entirely settled. (2) More directly, de Lille's mayoral term was marked by instability acute enough to impede stable governance. One should recall the protracted, public rupture with her own party that culminated in her departure in 2018 and the resignation of senior councillors in its wake (Davis, 2018). The instability of the Cape Town council over this period is also visible in the South African online news data. Referring back to Figure 3, Cape Town features far more prominently in instability-related coverage over 2010–2021 than from 2022 onward. The framing of GHL's term as stable coincides with the data, justifying treating the GHL administration as treatment rather than the 2011 attainment of a DA majority.

Why do we reject Dan Plato's return to office following de Lille's departure as our treatment onset? Firstly, Plato inherited a budget approved under another mayor's leadership. Secondly, given the disruption presented by key leadership's departure, we consider Plato's mayoral term as an interim restoration of balance. Lastly, while Plato's own departure from the DA in 2024 was not marked by the same acrimony as that of de Lille, it may at the very least be indicative

[10] It should also be conceded that Cape Town's high GPI is likely a result of structurally better service delivery prior to GHL's arrival in office. We argue that while service delivery may improve stability on the constituent front, it cannot shield the council from intra-party politics.

of some dissatisfaction with DA internal affairs. Regardless, the news data from Figure 3 shows that instability may only really have abated after the 2021 local elections.

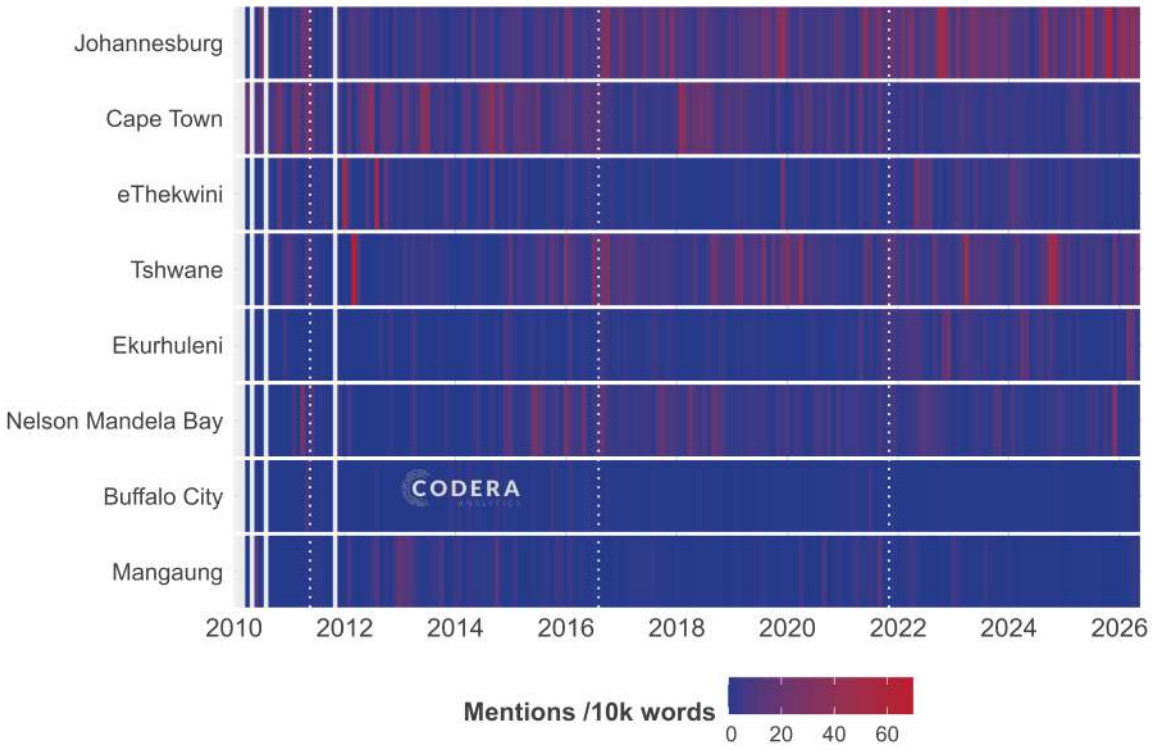

Figure 3: Monthly mentions of metros in coalition/ instability articles per 10,000 words.

Notes: Data sourced from BusinessLive, BusinessTech, Daily Maverick, Mail & Guardian, and SABC online archives. A regular expression algorithm identifies the number of articles per month containing terms related to metropolitan or local municipalities, coalitions, and council instability; metro name mentions are then aggregated. Some false positives are likely to persist: city names are frequently used in article headers or footers to denote the point of correspondence rather than the subject of the article. We believe these represent a negligible proportion of total metro name mentions.

Third, our design is potentially exposed to reverse causality. We treat governance stability as a cause of labour-market performance, but the relationship may run the other way or be self-reinforcing. A substantial economic-voting literature finds that voters reward incumbents for favourable economic conditions and punish them for adverse ones, with employment among the more consistent predictors of incumbent support (Kramer, 1971; Fair, 1978). GHL's tenure may thus have been stable in part because Cape Town enjoyed stronger economic conditions than the rest of the country, just as instability elsewhere may have followed from a deteriorating situation; indeed, much of the frustration with the ANC in other metros has been economic in origin.

Unfortunately, we cannot address this flaw within our research design. We would argue, however, that South African voting behaviour prioritises attributes extending beyond employment or economic output—race, service delivery, and spatial equity among them (Ferree, 2006; de Kadt and Lieberman, 2017). On this view, the economic-voting mechanism, largely estimated on United States elections, is weaker in the South African setting than the literature would imply. Our earlier argument is again relevant: that the DA held an outright majority in Cape Town from 2011 yet still suffered intra-party conflict suggests that stability is not governed by economic performance alone. This, in turn, points to a relative weakness in the reverse channel.

Fourth, the treatment window opens in the immediate aftermath of the COVID-19 pandemic,

and Cape Town's labour market may have rebounded more sharply than its peers for reasons unrelated to governance. The city's economy is comparatively exposed to tourism and hospitality, so part of the post-2022 absorption gain could simply reflect a sectoral recovery rather than the GHL administration. We cannot fully separate the two, but two features of the design mitigate the concern: our counterfactual Cape Town is constructed from other metros that experienced the same national pandemic shock, so a common recovery is differenced out, and we trim the most acute pandemic quarters in a robustness check (Section 4) to confirm the estimate does not rest on the recovery period alone.[11]

Fifth, besides the pandemic, the South African economy was further burdened by a deepening electricity crisis, with Eskom's load-shedding intensifying through the early 2020s. The City of Cape Town was partially insulated from this shock. The inherited Steenbras pumped-storage scheme—the only municipally operated plant of its kind in the country—enabled the City to shield its customers from up to two stages of load-shedding relative to the rest of the country during the worst load-shedding years. Because these years coincide with our treatment onset, our estimate may be inflated.

Lastly, the estimate may be confounded by Cape Town's pronounced in-migration. As noted above, using the absorption rate rather than employment levels removes the purely mechanical effect of a growing population, since the measure is a ratio. One should, however, remain sceptical of the composition of that inflow: if those most able to relocate are also those most able to find work, semigration could raise absorption through selection rather than through any improvement in local conditions. To the extent that migration occurs independently from the governance channel, we regard our estimate as an upper bound. That said, for employed migrants to be absorbed at all, the Cape Town economy must generate the jobs that sustain them, and the capacity to do so is plausibly itself a function of the stable governance whose effect we seek to estimate.

While the number of potentially confounding factors may seem alarming, it should be noted that this research is an initial attempt at answering a question that is becoming increasingly important in the political economy, both locally and abroad. Given the difficulty posed by the setting and by the limited data available to us, we hope that this design may open new pathways to understanding the role of local governance in promoting positive economic outcomes. Next, we proceed to discuss the results.

3 Results

Figure 4 visualises our estimation results for three estimation methods: two-way fixed effects difference-in-differences (TWFE DiD), synthetic control methods (SCM), and synthetic difference-in-differences (our preferred estimation).[12] Positively, the pre-treatment period in the DiD panel shows that the parallel trends assumption appears to hold visually. Additionally, and equally reassuring, is the close similarity between the shapes of the DiD and SDID counterfactuals. The SCM counterfactual, on the other hand, does not seem a good predictor of Cape Town's pre-treatment trend, owing to its inability to make up for level differences between Cape Town and other metros. The intercept term in the SDID estimation, which allows for a structural level difference between Cape Town and its donors, is what leads us to prefer

[11] It should be noted that cities like Johannesburg and eThekwini are also largely exposed to tourism owing to their respective proximity to safari, coastal, and resort destinations. This exposure is accounted for in the construction of the SDID counterfactual.

[12] The figure and related point estimates, as summarised in Table 1, were calculated using the `synthdid` R package, developed by Arkhangelsky et al. (2021).

this method.

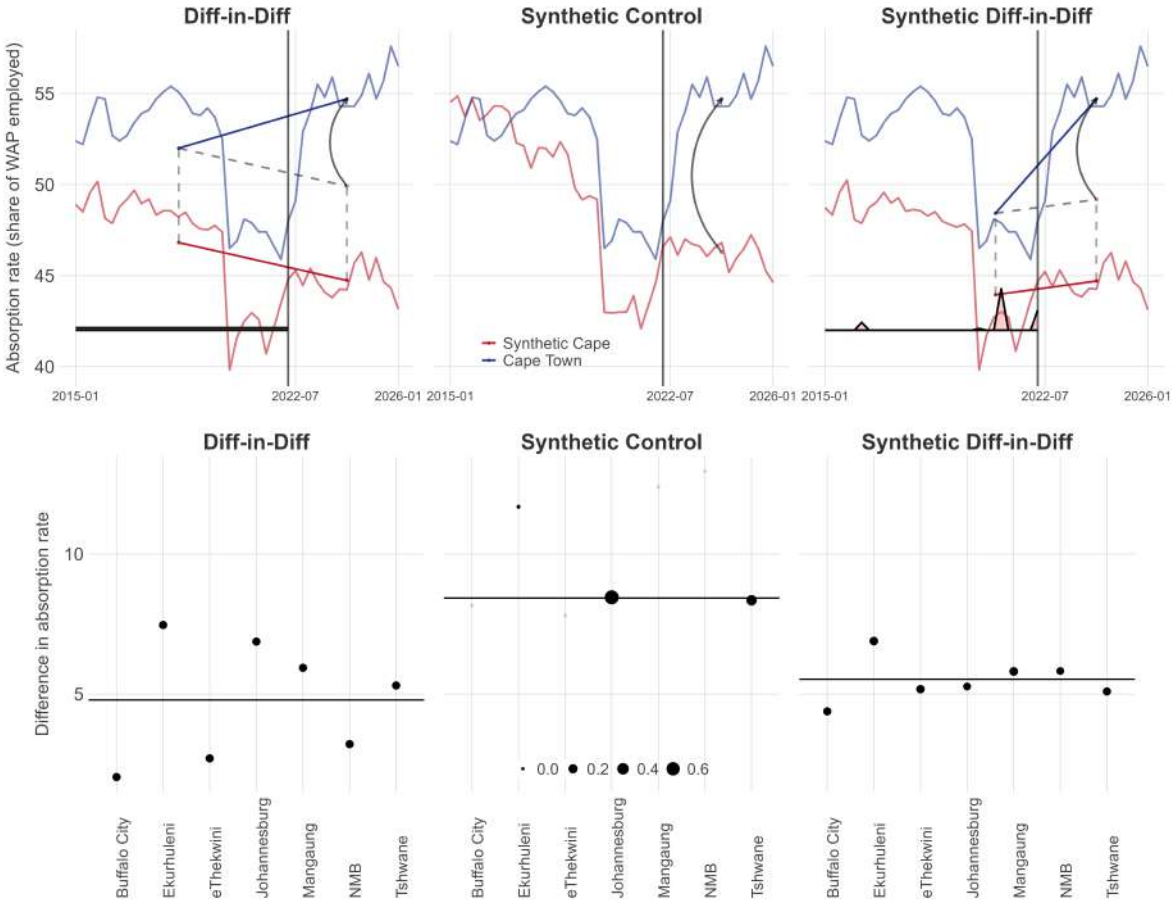

Figure 4: DID, SCM, and SDID estimates for Cape Town vs. a counterfactual Cape Town.

Notes: Top panel: Absorption rate trends of Cape Town and the weighted average of its metro counterparts. Time weights are reported below the trends. Bottom panel: Unit-level outcome differences, scaled by weight. Dot size denotes each donor's weight and the estimated effect is drawn by the horizontal line. Metros assigned zero weights are indicated by the $\times$ symbol.

Column 1 in Table 1 reports the ATTs of our respective estimators. On average over the post-treatment period, Cape Town's absorption rate was 5.53 percentage points higher than its synthetic counterfactual. This figure can be read as the employment absorption Cape Town would have sacrificed had it experienced the same council instability and mayoral turnover as other South African metros. Another way to interpret our results is that the stable, single-party mayoral tenure in Cape Town accounts for nearly half of the raw difference between the City and a weighted average of seven other South African metros, with the remainder reflecting pre-existing structural differences that predate the treatment period.

The credibility of the SDID estimate is supported by how closely the synthetic control tracks Cape Town's absorption rate prior to July 2022. The root mean squared prediction error (RMSPE) over the pre-treatment period is 1.39pp—roughly 2.7 percent of Cape Town's mean pre-treatment absorption rate—indicating a good pre-treatment fit. By contrast, the SCM estimate overstates the effect because, lacking an intercept term, it cannot absorb the structural level difference between Cape Town and its donors. Its counterfactual yields a substantially higher pre-treatment RMSPE of 3.1pp, confirming that it fails to reproduce Cape Town's pre-treatment trajectory. Our DiD estimate—a relatively conservative ATT of 4.79 percentage points—is more reassuring, again lending credibility to SDID and its level-matching opera-

tion.[13]

Table 1: Comparison of Treatment Effect Estimates: DID, SCM, and SDID

	(1) Baseline	(2) COVID-trimmed	(3) Onset July 2023
Panel A: Treatment effect (ATT, pp)			
DID	4.789	4.849	5.201
SCM	8.427	8.712	9.258
SDID	5.526	5.679	2.313
Panel B: Pre-treatment fit (RMSPE, pp)			
DID	1.27	1.36	1.74
SCM	3.10	2.78	3.71
SDID	1.39	1.49	3.09
Panel C: Support			
T_0 (pre-periods)	29	24	33
T_0 effective (SDID)	2.31	1.93	2.46
T_1 (post-periods)	15	15	11

Notes: Column (1) is the baseline (onset July 2022). Column (2) drops the five deepest pandemic quarters (2020–Q2 through 2021–Q2). Column (3) moves onset to July 2023, discarding the administration's first full budget year; the later onset forces the counterfactual to fit four additional, post-2022 quarters in which Cape Town had already begun diverging, raising pre-treatment RMSPE; T_0 effective is the inverse-Herfindahl of the SDID time weights. Inference is reported separately via permutation tests in Section 4. Data is missing across all units for 2021–Q4.

The bottom panel of Figure 4 shows the distribution of weights assigned to units in our control group. Donors in our SDID counterfactual are very similarly weighted, with all metros contributing 11–17 percent to the final control. Time weights, on the other hand, are assigned much more sharply, with 2021–Q1 receiving 58.3 percent of the importance. 2022–Q2 and 2016–Q2 take 28.2 and 11.1 percent, respectively. Mechanically, this concentration is simply how the algorithm minimises the pre-treatment RMSPE, and a somewhat comforting check is the similarity between our SDID and DiD results, given that the latter assigns equal weights to all donor units and time periods. The concentration does, however, carry a substantive concern: with the two most dominant donor periods falling in the COVID-19 recovery years, it raises the question of how this period could contaminate our results. We address this directly by re-estimating on a panel that drops the deepest pandemic quarters (2020–Q2 through 2021–Q2); the estimate is essentially unchanged at 5.68pp (column 2 of Table 1), indicating the result does not hinge on the COVID-recovery window, though the effective number of pre-periods falls further (to 1.93), reflecting the reduced pre-treatment support.

A related but distinct COVID concern presents itself in what we assume would be explained by the inbound travel industry. In the quarters immediately preceding treatment, Cape Town's absorption rate falls below its synthetic counterfactual (gaps of -2.18pp and -1.18pp at 2022–Q1 and 2022–Q2), as the inland donor metros recovered from the COVID-19 pandemic ear-

[13] DiD's pre-treatment RMSPE is 1.27pp, marginally below SDID's despite being the more constrained model. SDID spends some of its pre-treatment fit on trajectory-matching through its unit and time weights, whereas DiD imposes equal weights throughout, so this result is unsurprising. The proximity of the two models' ATTs (4.79 and 5.53pp) and their near-identical pre-treatment fit is reassuring of the model's performance.

lier than Cape Town's more tourism-dependent economy.[14] A concern is that a portion of the post-treatment gap may reflect Cape Town completing a delayed recovery rather than a treatment effect, which would inflate our baseline estimate. We test this directly by moving treatment onset one financial year later, from July 2022 to July 2023, which pushes the onset beyond the recovery window and trims the post-treatment period. The estimate, reported in column 3 of Table 1, falls to 2.31pp. The decline is consistent with the upward bias that a tourism rebound would introduce into the baseline. On the other hand, by dropping GHL's first budget we are arguably discounting the administration's most impactful policies and forfeiting over 18 months in office. Furthermore, one should not ignore that stable governance—and the policies it enables leadership to push, such as an emphasis on safety—may itself be conducive to greater inbound demand. Altogether, these considerations lead us to treat 2.31pp as a conservative lower bound rather than a point estimate. Our DiD and SCM estimates remain near their baselines at 5.2pp and 9.26pp, as neither re-weights pre-treatment periods and so neither is sensitive to the shifted onset. We subject the estimate to a full suite of placebo tests in the next section.

4 Placebo Tests

The existing synthetic control tradition (Abadie and Gardeazabal, 2003; Abadie et al., 2010, 2015) provides a valuable toolkit of robustness checks we can perform to confirm the validity of our estimates. Andersson (2019), using SCM to study the impact of a 1990 carbon tax and VAT hike on carbon emissions, also guides the following exercises. Beyond the initial checks done in the results section,[15] we perform the following placebo tests: "in-time," "in-space," and "leave-one-out." Finally, we perform the same study using full-time employment data from the South African Revenue Service (SARS).

The in-space test drops Cape Town from the panel and runs through the remaining donor units, iteratively assigning treatment to each and running the estimation. Essentially, we get point estimates, RMSPEs, and draw gap plots for each of the donors, based on a unique counterfactual of $N = 6$ other metros. Panel A of Figure 5 plots the gap measure for each in-space permutation, with Cape Town's original highlighted in blue. Cape Town's pre-treatment gap is centred around zero and diverges dramatically in the post-treatment period, relative to other metros. Additionally, Panel B shows each metro's post- to pre-treatment MSPE ratio—Cape Town clearly dominating. While classical standard error inference is not possible when using SDID with a single treatment unit, we obtain a pseudo p-value of $p = 0.125$.[16]

Arkhangelsky et al. (2021) propose a placebo variance estimation algorithm that employs the iterative sampling logic of the in-space test. Using their algorithm, we collect the placebo estimates $\hat{\tau}^{(b)}$ and define placebo variance as follows:

$$\hat{V}_{\tau}^{\text{placebo}} = \frac{1}{B} \sum_{b=1}^{B} \left(\hat{\tau}^{(b)} - \bar{\hat{\tau}} \right)^2, \tag{2}$$

[14] "Synthetic counterfactual" here refers to the intercept-adjusted SDID counterfactual, which shifts the weighted donor average upward by the pre-treatment level offset and is plotted as a dotted line in Figure 4. Cape Town's raw absorption rate remains above the unadjusted donor average throughout the sample; the negative pre-treatment gaps are measured against the level-adjusted counterfactual the estimator targets.

[15] Dropping the periods most contaminated by COVID-19 variation and shifting treatment onset one year later to account for tourism rebound.

[16] Pseudo p-values are calculated by measuring the fraction of metros with point-estimates at least as large as our treated unit, Cape Town. Since Cape Town itself is the largest in our setting of $N = 8$ total metros, $p = 0.125$ is the minimum, and Cape Town's result remains the most significant.

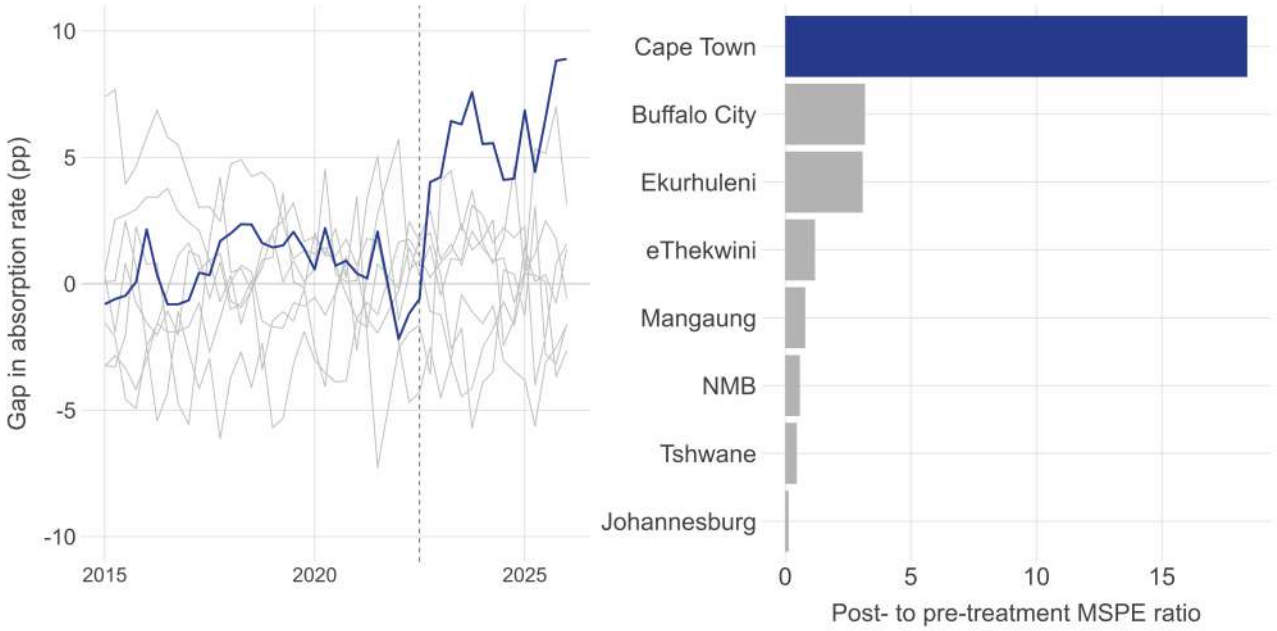

Figure 5: In-space placebo results.

Notes: Gap measures the difference between "treated" unit and the intercept-adjusted SDID counterfactual.

where $\bar{\hat{\tau}} = \frac{1}{B}\sum_{b=1}^{B}\hat{\tau}^{(b)}$. From here, calculating the placebo standard error is trivial. For these standard errors to be valid, however, we require homoskedastic errors across units. As this assumption cannot be tested with a single treated unit, we rely instead on the first-differenced pre-treatment volatility as a proxy. Cape Town's outcome is the least volatile, with the standard deviation of its first differences equating to 1.43pp, whereas donors range between 1.55 and 3.17pp. The placebo draws are therefore constructed from units noisier than our treatment unit, which is more likely to place upward bias on $\hat{V}_{\tau}^{\text{placebo}}$ and thereby renders our inference conservative rather than fragile. Recalling the implied point estimate of 5.5pp, we compute a placebo standard error of 1.76, and the resulting confidence interval $(2.08, 8.97)$ excludes zero at the 5% level even under this conservative variance.

We test for temporal robustness by repeatedly shifting the treatment onset backwards in time. Owing to the seasonality present in employment data, we keep quarters constant while running the estimations for one year prior.[17] It should also be noted that observations from July 2022 onwards are cut from the sample to avoid contaminating our placebo with time periods in the real post-treatment period. Figure 6 reports the results for four separate in-time tests, none of which shows any significant shift in absorption rates.

Setting treatment onset to 2021–Q3 and 2020–Q3 yields ATTs of -1.9 and -1.5, respectively. These observations mainly capture Cape Town's sharp downturn owing to the COVID-19 crisis and may, at least suggestively, lend credibility to our interpretation of the estimation in column 3 of Table 1 representing a conservative lower bound of the true tenure stability effect (given that these downturns may capture the loss of tourism and other heavily exposed jobs during the crisis period). Effectively zero estimates in Panel C and Panel D show that our preferred estimate is robust to changes in treatment onset.

One should ask whether our counterfactual Cape Town is robust even when we drop one of the sample of potential donors; i.e., "leaving-one-out". The results show that our estimate re-

[17] Another rationale is that we want to test the onset of prior budgets.

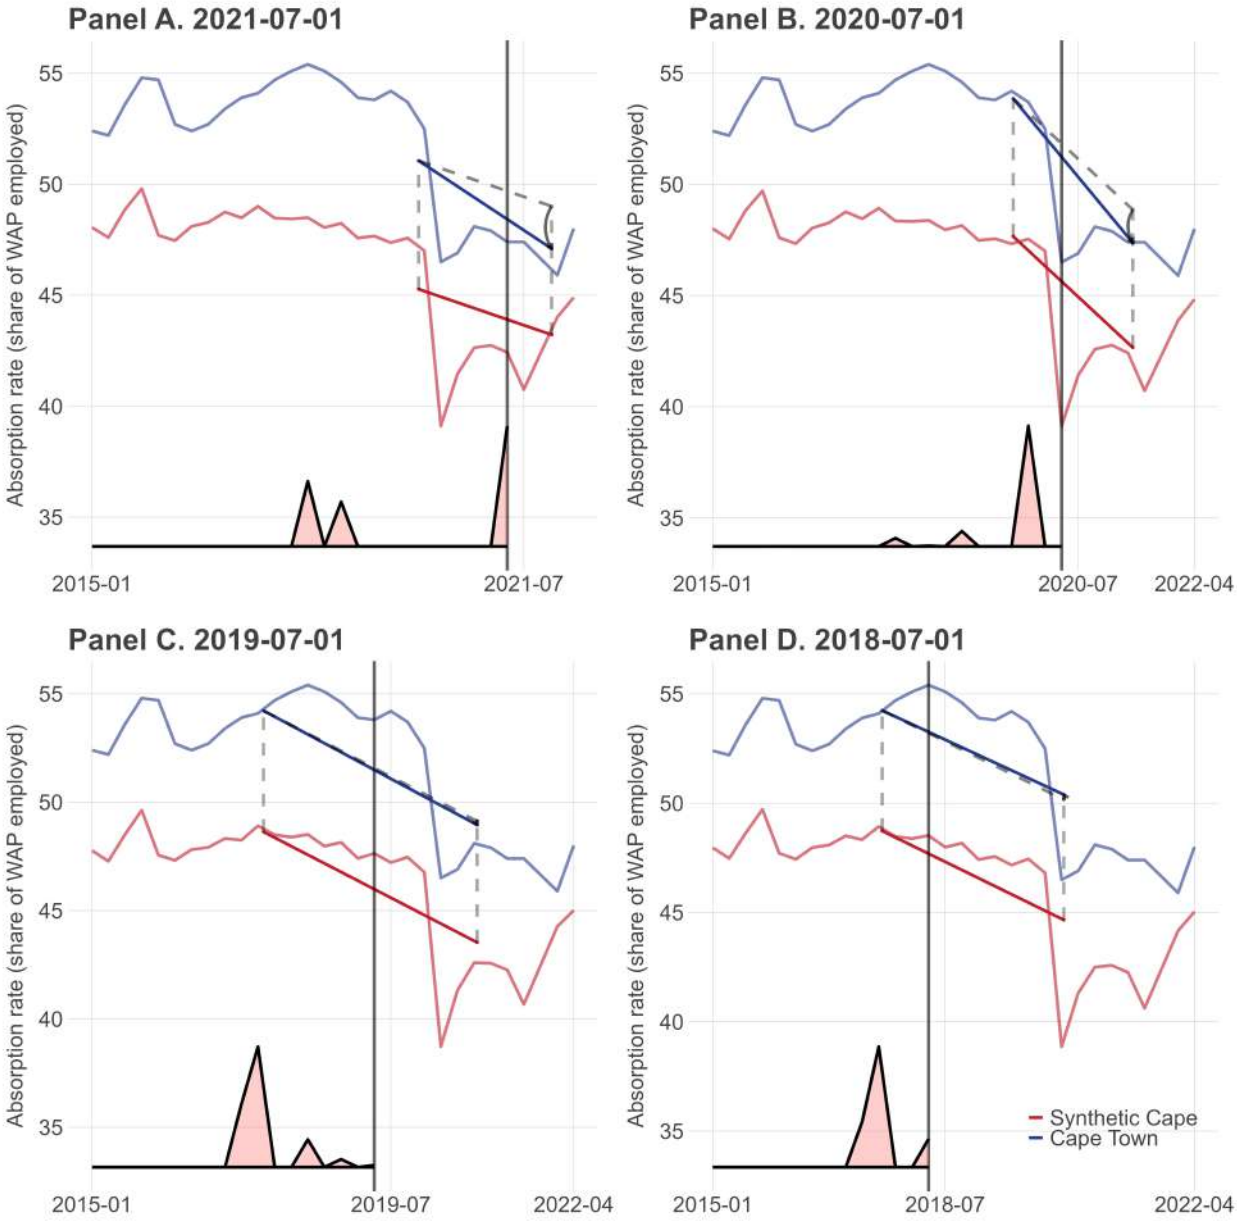

Figure 6: In-time placebo results.

mains relatively unchanged. Dropping Ekurhuleni causes the largest downward shift resulting in an estimated ATT of 4.4pp across the remaining $N = 6$ donors. Leaving out Nelson Mandela Bay, conversely, raises the estimated effect to 5.68pp. Remaining placebo estimates remain within this range. As these results and their visualisation in Figure 7 clearly show, our main results remain robust to the dropping of single units from our donor pool, regardless of the donor.

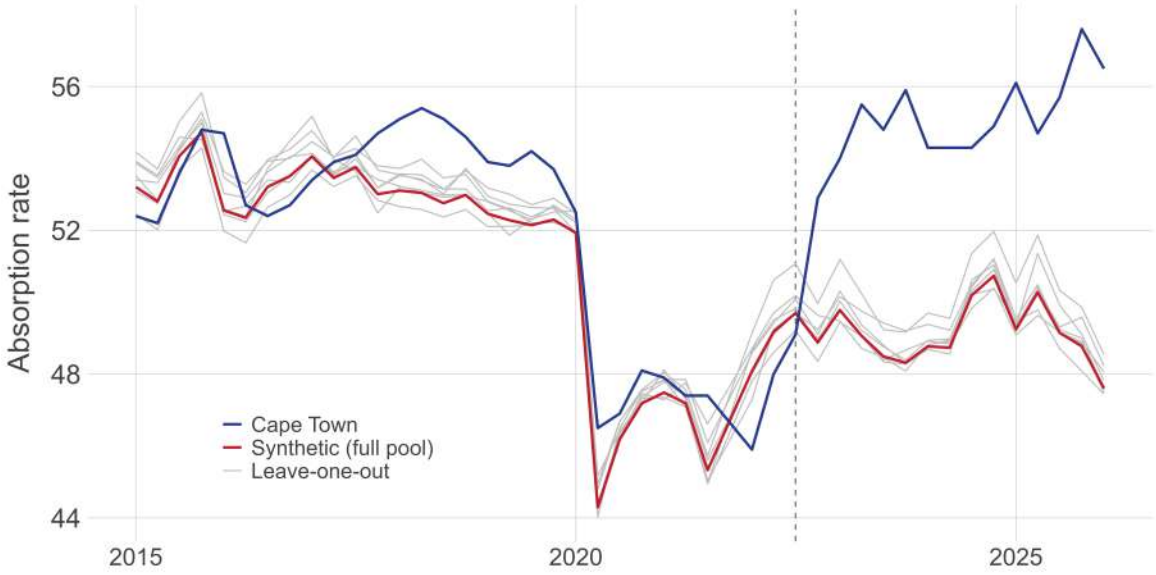

Figure 7: Leave-one-out results.

Notes: Counterfactuals are raised to the same intercept as Cape Town for clearer exposition.

Finally, we rerun our main estimators and placebo tests using full-time-equivalent employment from the Spatial Tax Panel. Results are summarised in Table 2 and the related figures can be found in our supplementary appendix. The STP-FTE outcome corroborates the absorption-rate result on the sign and timing of the effect while the strength of the effect is attenuated. The baseline SDID estimate of 0.038 log points—a 3.9 percent rise in formal full-time-equivalent employment—is robust across specifications, rising to 0.042 when the deepest pandemic months are dropped and to 0.053 under a July 2023 onset. Notably, and unlike the QLFS result, the later onset raises the estimate, implying that the formal-employment response to stable tenure was delayed. Given that policies require time to take effect, this finding is entirely sensible. The result may also point to differences in the recovery dynamics of formal and informal employment in the wake of the pandemic. The monthly frequency provides a healthier pre-treatment support than the quarterly QLFS panel (T_0 effective of 4.16 against 2.31), and all three estimators agree closely and fit comparably, so the result is not an artefact of estimator choice.[18]

The placebo evidence, however, is weaker for the STP-FTE data. Cape Town remains rank one of eight in every specification, but the margin is small. The post/pre MSPE ratios of the in-space placebos are far more compressed than in the QLFS case (likely a consequence of working in log units) with Mangaung trailing Cape Town closely. Similarly, in the gap plot (Figure A.3), Cape Town's post-treatment divergence is less pronounced. The in-time placebos are likewise less conclusive, with no sharp null at the false onset. Nevertheless, it is encouraging that the estimated effects (0.016, 0.015, -0.010, 0.021) are all smaller than the 0.038 base-

[18] Unlike the absorption-rate specification, SCM fits the logged FTE series nearly as well as SDID, since logging compresses the metro-level differences that the SDID intercept otherwise absorbs.

Table 2: Comparison of Treatment Effect Estimates (STP FTE Data)

	(1) Baseline	(2) COVID-trimmed	(3) Onset July 2023
Panel A: Treatment effect (ATT, log points)			
DID	0.067	0.070	0.073
SCM	0.050	0.051	0.053
SDID	0.038	0.042	0.053
Panel B: Pre-treatment fit (RMSPE, log points)			
DID	0.0327	0.0346	0.0345
SCM	0.0198	0.0208	0.0225
SDID	0.0212	0.0190	0.0186
Panel C: Support			
T_0 (pre-periods)	112	97	124
T_0 effective (SDID)	4.16	3.64	2.30
N_0 effective (SDID)	2.04	1.90	1.77
T_1 (post-periods)	32	32	20

Notes: Outcome is log formal full-time-equivalent (FTE) employment from the SARS Spatial Tax Panel.

line. Consequently, we read the STP-FTE result as independent administrative-data support for the direction of the stable-governance effect, while conceding that its synthetic control depends heavily on a single donor (Tshwane assigned a weight of 0.67–0.71) and that its placebo separation is fairly narrow.[19]

5 Industry Decomposition

A powerful feature of the Spatial Tax Panel is the availability of full-time-equivalent employment data at the industry level.[20] Using the STP-FTE data, we decompose changes in full-time employment by industry in order to show how stable governance heterogeneously affects different parts of the economy. Furthermore, such an exercise helps us understand whether the estimated effect reflects our treatment, or whether it is instead attributable to industry-specific confounders we previously noted.

In results tabulated in Appendix Table A.1, we run the SDID estimator separately across each of the SIC-1 industry classifications. Each industry is thus assigned its own synthetic Cape Town, with donor weights free to differ across industries so as to best reproduce each industry's pre-treatment path. This sharpens the per-industry counterfactual, though at the cost of noisier fits in thin sectors where Cape Town's employment base is small.

For each industry j we report the SDID estimate $\hat{\tau}_j$ alongside two descriptive statistics. Observed growth is the raw change between the pre- and post-treatment window means, $g_j = (\text{FTE}_j^{\text{post}}/\text{FTE}_j^{\text{pre}} - 1) \times 100$, where $\text{FTE}_j^{\text{pre}}$ is mean monthly FTE employment in industry j over

[19] On the weights, our leave-one-out placebo for the STP-FTE data reveals that a removal of Tshwane from the donor pool does not change the results. Therefore, we argue this weight is not a cause of concern.

[20] At the lowest level, data is available by 5-digit industry classification, for which 521 unique industry codes exist.

the 12 months preceding treatment and $\text{FTE}_j^{\text{post}}$ its analogue over the final 12 months of the sample. Weighting each industry's implied proportional effect by its pre-treatment employment share gives a bottom-up counterpart to our headline estimate,

$$\hat{\tau}^{\text{BU}} = \sum_j \frac{\text{FTE}_j^{\text{pre}}}{\sum_k \text{FTE}_k^{\text{pre}}} \left(\exp(\hat{\tau}_j) - 1\right), \tag{3}$$

which implies that stable administration may have raised total FTE employment in Cape Town by roughly 3 percent between July 2022 and February 2025. The proximity to our 3.9 percent aggregate estimate is reassuring and the two estimates need not coincide: donor weights differ across industries and the log transformation does not commute with aggregation. As such, $\hat{\tau}^{\text{BU}}$ is better read as a consistency check than a competing estimate.

Finally, we translate each estimate into an implied level effect,

$$\widehat{\Delta L}_j = \text{FTE}_j^{\text{pre}} \left(\exp(\hat{\tau}_j) - 1\right), \tag{4}$$

which represents the number of FTE positions attributable to treatment in industry j, evaluated at its pre-treatment base. Summing across industries suggests that stable governance over the GHL tenure may have supported nearly 37,000 FTE positions relative to a counterfactual composed of the donor metros.

For expositional simplicity, and to uncover the broad mechanisms underlying Cape Town's employment changes, we split the 21 industry classifications into seven groups and repeat the procedure. Table 3 reports the group-level SDID estimates alongside observed characteristics. We focus in particular on tourism, since the post-COVID recovery window presents a substantial threat to our identification strategy. The tourism-exposed accommodation and food services industry records the largest raw growth of any group over the post-treatment period, at 28.4 percent; however, this is largely the data revealing the expected. Once Cape Town's tourism growth is measured against its synthetic counterpart—now weighted toward metros that best reproduce Cape Town's own pre-treatment tourism path—the estimated excess is far more modest. This finding considerably abates our earlier concern that the results may be confounded with a larger than expected tourism rebound in Cape Town. After differencing out common tourism and COVID recovery trends, the remaining residual of 4.5 percent is the only growth plausibly attributable to Cape Town's local environment.

As a brief illustration of how SDID weights the donor pool, the tourism counterfactual loads on tourism-heavy metros: eThekwini (0.27) and Tshwane (0.20) take the largest shares, corroborating our earlier assertion that their proximity to beach and bushveld getaways—as well as other amenities like golf or luxury resorts—makes them comparable inbound-travel markets. Furthermore, Mangaung (0.19) may capture a domestic accommodation rebound as a popular halfway house between Cape Town and Gauteng, whereas Johannesburg (0.18) enters as a cultural and dining hub. Together, these four metros represent 84 percent of the synthetic control for Cape Town's accommodation and food services. The logic of the weighting is reinforced by Buffalo City receiving a zero weight, consistent with it being among the least tourism-oriented metros in the donor pool.

Although the tourism- and migration-exposed groups exhibit the largest raw growth rates and treatment effects, their small employment shares limit their contribution to the aggregate: together they account for roughly 4,700 of the implied FTE positions, against a net total of 24,348 implied by the industry group estimations.[21] Even under a strong assumption that

[21] The grouped total is smaller than the industry-level sum due to aggregation of logs and divergence in weights

Table 3: Grouped Sector SDID Decomposition

Sector Group	SIC Codes	Share	Growth	$\hat{\tau}_j$	$\widehat{\Delta L}_j$
Tourism exposed	9	0.044	28.4%	0.044	2,426
Tourism & migration adj.	12, 18	0.024	12.2%	0.073	2,276
Business & services	10, 11, 13, 14	0.260	8.1%	0.034	11,090
Tradables & construction	3, 6, 7, 8	0.363	1.9%	−0.002	−806
Public admin	15	0.129	−0.6%	0.018	2,804
Health and Education	16, 17	0.102	1.7%	0.029	3,703
Other	1, 2, 4, 5, 19, 20, 21	0.077	12.2%	0.029	2,855
Total		1.000	—	—	24,348

Notes: Share is the pre-treatment employment share. Growth is the observed growth between the mean pre-treatment and final-sample year windows, not a treatment effect. $\hat{\tau}_j$ is the sectoral SDID estimate in log points; inference is unreliable with seven donors, so SEs are omitted and estimates are descriptive. $\widehat{\Delta L}_j$ is the implied level effect from Eq. 4.

attributes all of tourism and migration's measured effect to confounding rebounds, the headline estimate would fall only modestly. We therefore contend that any over-attribution to tourism and semigration is minor in magnitude.

The tradables and construction group offers a similar insight in the opposite direction. Despite Cape Town having been somewhat insulated from load-shedding during its worst years in 2022 and 2023—which might lead one to expect larger relative gains—this energy-intensive group records a small decline in FTE employment relative to its synthetic counterpart. We argue that metros which suffered more severe load-shedding had more suppressed industrial activity to recover, so their sharper post-crisis rebound raises the counterfactual against which Cape Town is measured. Cape Town's flat tradables performance thus possibly reflects the absence of a strong rebound, likely attenuating our headline estimate.[22]

Cape Town's FTE employment growth is, therefore, concentrated in financial, ICT, and administrative support sectors. We identify three plausible channels: (1) stable governance may sustain the policies and conditions that help firms expand; (2) service-delivery differentials—strong provision in Cape Town, or its deterioration elsewhere—may draw firms and skilled workers to semigrate; and (3) large multinational investments, several of which predate the GHL administration, expand employment opportunities in the metro. We cannot fully separate these channels, and instead reason through them informally in our discussion section.

6 Discussion

This research, to our knowledge the first of its kind in a South African setting, considers an important question: do stable, majority-run governments perform better than coalitions or council leadership facing political pressure? Our synthetic difference-in-differences approach finds that the stable tenure of the Hill-Lewis administration contributed 2.3–5.5pp to the City

used to construct counterfactuals. What is interesting is that our overall estimates differ non-monotonically as we reduce the level of aggregation. The seven-sector industry decomposition produces the smallest overall estimated effect of 2 percent.

[22] This load-shedding account is offered conjecturally. Appendix Table A.1 shows that relative job losses in manufacturing and retail were largely offset by gains in construction; since building activity is a forward-looking indicator of economic confidence, the group's post-treatment performance is arguably stronger than the headline decline suggests.

of Cape Town's absorption rate in excess of a counterfactual constructed from the remaining seven South African metros.[23] Furthermore, the broad agreement between our QLFS-based absorption measure and the administrative formal-employment panel lends extra support to the direction of our estimates.

To what extent do majority governments have more flexibility to enact policy and thereby promote economic growth and job creation? Hung councils are frequently plagued by infighting and political competition, potentially sabotaging an incumbent mayor's plans in order to deflect blame and reap political profits at the next round of elections. We therefore propose that the mechanism by which stable majority governments foster positive economic outcomes is their ability to act decisively and enact policy or perform regular duties with less resistance.[24]

We also ask how much of Cape Town's employment growth reflects the post-COVID recovery of its tourism sector rather than governance. The two are difficult to disentangle, and we do not claim to fully separate them. Using the STP-FTE panel, however, we conclude that Cape Town's excess growth in tourism-exposed industries is not the main driver behind its employment figures outpacing those of its contemporaries. Additionally, it is entirely plausible that Cape Town's focus on tourism and on promoting itself on the global stage facilitated a larger recovery in the inbound industry than would otherwise have occurred. In such a case the tourism rebound is partly a channel of governance rather than a confound external to it. Policies consistent with this point of view include the City's deployment and expansion of dedicated tourism-safety units from late 2022 (Western Cape Government, 2022).[25]

An important caveat to the employment figures is that we remain uncertain about the distribution of these jobs along the lines of income groups. Aggregate absorption reveals little about employment equity, and there is growing concern that the City has under-prioritised spatial redistribution and left some of its most vulnerable residents behind. While absorption and FTE may have risen, one should ask whether the benefit accrues to long-standing residents or disproportionately to those who have since migrated into the city. Future research could benefit from decomposition of employment into income brackets to partially answer these questions; though tying such descriptive results to the causal effect of stable tenure will remain difficult. A related concern is the rise in rates and tariffs accompanying buoyant property prices. While these plausibly reflect the attraction of migrants owing to better service delivery or economic opportunities, we make no claim that the City's fiscal choices are optimal, and an assessment of their incidence is beyond the scope of this paper.

Next, we discuss the three channels identified in our industry decomposition that may have reinforced stronger-than-counterfactual growth in Cape Town's highly skilled FTE employment. Stable governance, channel (1), is our target estimand and plausibly makes up the largest proportion of our estimate. The semigration channel (2) remains ambiguous; if service delivery is itself a product of stable governance, then the deterioration of municipal provision in unstable metros and the resulting exodus to Cape Town means that semigration is a supportive mechanism. Nevertheless, given the violation of no-interference, it remains upward biasing. Finally, channel (3): provided Cape Town attracted disproportionate multinational commitment prior to our treatment, such investment likely inflates our estimates. Several large commitments—Amazon Web Services' Cape Town region among them—were made before Hill-Lewis took office. The employment generated is realised with a lag and accrues across

[23] We stress that this is not a measure of partisan effect; instead, it reflects the benefits of political stability and continuity.

[24] It should be conceded that our current study cannot test this directly.

[25] The City's unit predates our treatment, yet expanded under the Hill-Lewis administration; we therefore describe an intensification of policy.

periods extending into our post-treatment window irrespective of who governs. Investment inflows, supported by inherited policy and a higher baseline stability, motivate our framing of the skilled-employment and headline estimates as an upper bound.

We offer a final conjecture: stable governance may itself underpin these confounders. Consider initiatives such as Cape Town Air Access and the "Silicon Cape" push, which originated prior to Hill-Lewis's accession to the mayoralty. We hold that the contribution of governance stability may lie less in initiating such policies than in allowing them to run uninterrupted long enough to be refined and to compound. Continuity of stable, single-party rule is conducive to that persistence: policy programmes are less exposed to the disruption, renegotiation, or abandonment that mid-term turnover and coalition instability can bring. What our estimates modestly suggest, then, is that administrative stability may reinforce the attraction of firms and professionals, with the potential for agglomeration effects under favourable conditions.

This research also casts light on how political organisations may frame their governance record during electoral periods—a point of particular relevance ahead of the local elections on 4 November 2026. Although the gap between Cape Town's absorption rate and that of other metros is large and widening, our results show that not all of it is attributable to the Hill-Lewis administration. Cape Town begins from a higher employment base and has remained under DA majority since 2011; much of the raw gap is therefore structural and predates the treatment period. Our headline estimate suggests that roughly half of the raw gap between Cape Town and the donor metros emerged during the post-treatment period; we associate this emergent component with the stable-governance effect, while emphasising that it remains entangled with the confounders discussed above.

Further research could leverage the SARS Spatial Tax Panel at the municipal level to compare hung councils with those under majority rule within provinces, holding constant much of the regional heterogeneity that separates treatment and control units. Crucially, such a design would provide multiple treated municipalities with staggered onsets, allowing for valid statistical inference of the kind our single-treated-unit setting cannot support, thereby moving the stable-governance question from suggestive estimation toward credible identification. Though, such research would require tremendous effort in the construction of an accurate mayoral handover database.

Considering the factors that may weaken identification, alongside the limitations of the data, these results must be taken as a first-attempt estimate of a stable-governance effect in South Africa. The purpose of this paper is to propose an avenue for research to evaluate the performance of local government in South Africa and guide the current policy debate regarding coalition governments across the spheres of South African government.

References

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

Abadie, A. and Gardeazabal, J. (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1):113–132.

Alesina, A. and Drazen, A. (1991). Why are Stabilizations Delayed? *The American Economic Review*, 81(5):1170–1188.

Andersson, J. J. (2019). Carbon Taxes and CO2 Emissions: Sweden as a Case Study. *American Economic Journal: Economic Policy*, 11(4):1–30.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic Difference-in-Differences. *American Economic Review*, 111(12):4088–4118.

Bertrand, M. and Schoar, A. (2003). Managing with Style: The Effect of Managers on Firm Policies. *The Quarterly Journal of Economics*, 118(4):1169–1208.

Besley, T., Montalvo, J. G., and Reynal-Querol, M. (2011). Do Educated Leaders Matter? *The Economic Journal*, 121(554):F205–F227.

CoGTA (2024). Local Government: Municipal Structures Amendment Bill, 2024. Department of Cooperative Governance and Traditional Affairs, Government Gazette No. 50682, 21 May 2024.

Davis, R. (2018). Chaos at De Lille's 'Farewell' as Councillors Resign, Citing Party Racism. Daily Maverick.

Dayimani, M. (2022). Nelson Mandela Funeral Trial: Accused Who Was Paid R1.3m Claims He Did Not Know Money Was Stolen. News24.

de Kadt, D. and Lieberman, E. S. (2017). Nuanced Accountability: Voter Responses to Service Delivery in Southern Africa. *British Journal of Political Science*, 50(1):185–215.

de Vos, P. (2021). The Constitutional-Legal Dimensions of Coalition Politics and Government in South Africa. In Booysen, S., editor, *Marriages of Inconvenience: The Politics of Coalitions in South Africa*, chapter 9, pages 235–266. Mapungubwe Institute for Strategic Reflection (MISTRA), Johannesburg.

Fair, R. C. (1978). The Effect of Economic Events on Votes for President. *The Review of Economics and Statistics*, 60(2):159–173.

Ferree, K. E. (2006). Explaining South Africa's Racial Census. *Journal of Politics*, 68(4):803–815.

Ferreira, F. and Gyourko, J. (2009). Do Political Parties Matter? Evidence from U.S. Cities. *The Quarterly Journal of Economics*, 124(1):399–422.

Good Governance Africa (2021). Governance Performance Index: South Africa 2021. Governance insights and analytics report, Good Governance Africa, Johannesburg.

Good Governance Africa (2024). Governance Performance Index: South Africa 2024. Governance insights and analytics report, Good Governance Africa, Johannesburg.

Grootes, S. (2019a). The ANC's Next Big Test: eThekwini Mayor Zandile Gumede's Corruption Case. Daily Maverick.

Grootes, S. (2019b). With Zandile Gumede's Removal, Tremors of Change Reverberate from the KZN ANC. Daily Maverick.

IEC (2021). Municipal Election Results Dashboard. Electoral Commission of South Africa.

Jones, B. F. and Olken, B. A. (2005). Do Leaders Matter? National Leadership and Growth Since World War II. *The Quarterly Journal of Economics*, 120(3):835–864.

Kerr, A. and Wittenberg, M. (2019). Earnings and Employment Microdata in South Africa. WIDER Working Paper Series wp-2019-47, World Institute for Development Economic Research (UNU-WIDER).

Kramer, G. H. (1971). Short-Term Fluctuations in U.S. Voting Behavior, 1896–1964. *American Political Science Review*, 65(1):131–143.

Makhaye, C. (2025). Promises and Perils — The Rise and Fall of Durban's Post-Apartheid City Fathers. Daily Maverick.

Mongale, O. and Nkwatle, B. (2025). From Policy to Practice: Strengthening Local South African Government Coalitions through Legislation ahead of the 2026 Elections. Policy & Practice Brief 068, African Centre for the Constructive Resolution of Disputes (ACCORD), Durban.

National Treasury (2020). Mangaung Intervention in Terms of Section 139 (5) (a) & (c). Media Statement, Ministry of Finance, Republic of South Africa.

Nell, A. and Visagie, J. (2025a). Addressing the Potential for Head Office Bias. Technical note, Spatial Economic Activity Data (SEAD-SA), Spatial Tax Panel.

Nell, A. and Visagie, J. (2025b). Spatial Tax Panel 2014–2025.

Paret, M. and Runciman, C. (2023). Voting Decisions and Racialized Fluidity in South Africa's Metropolitan Municipalities. *African Affairs*, 122(487):269–298.

Roubini, N. and Sachs, J. D. (1989). Political and Economic Determinants of Budget Deficits in the Industrial Democracies. *European Economic Review*, 33(5):903–933.

Statistics South Africa (2026). Quarterly Labour Force Survey (QLFS). Statistical Release P0211. Pretoria: Statistics South Africa.

Tandwa, L. (2023). Buffalo City Mayor Recalled, Wants to Meet National ANC Officials. Mail & Guardian.

Tladi, M. (2024). Antagonistic Politics and Coalitions: South Africa's Local to National Landscape. *Politikon*, 51(3-4):143–162.

Tolsi, N. (2019). Zandile Gumede: The Iron Lady of Inanda. Mail & Guardian.

von Fintel, D. (2023). Monitoring South Africa's Metropolitan Economies: A Survey of the Data Landscape. ERSA Working Paper 13, Economic Research Southern Africa. Discussion Document 13.

Western Cape Government (2022). Western cape police oversight and community safety launches tourism safety units. South African Government, gov.za, 29 November 2022.

Yao, Y. and Zhang, M. (2015). Subnational Leaders and Economic Growth: Evidence from Chinese Cities. *Journal of Economic Growth*, 20(4):405–436.

A Supplementary results

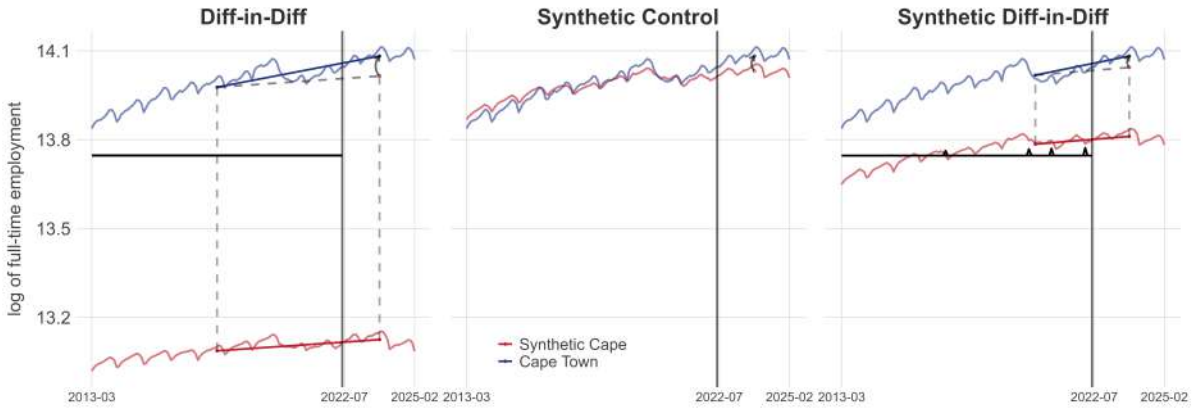

Figure A.1: DID, SCM, and SDID Estimates for Cape Town (STP-FTE Data)

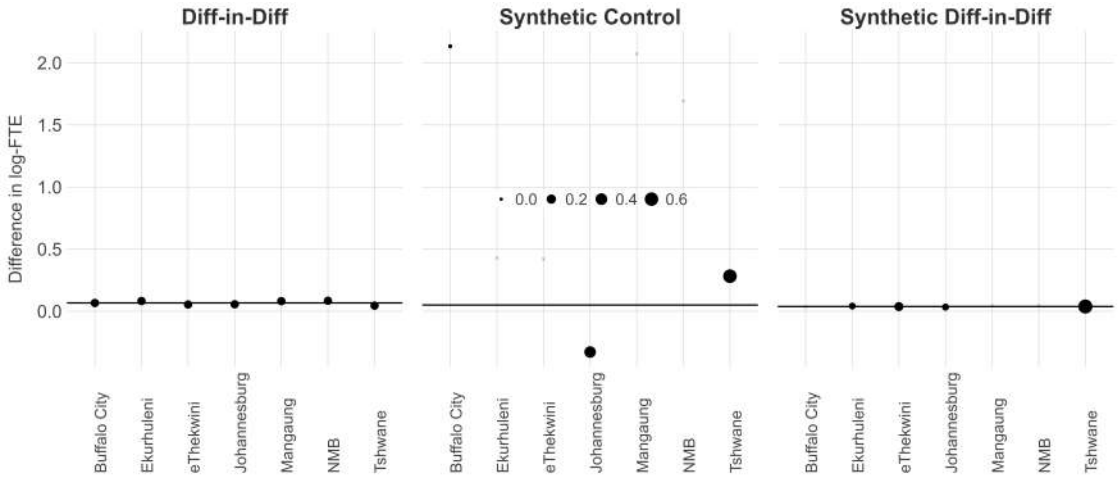

Figure A.2: Control Weights of DID, SCM, and SDID Estimators (STP-FTE Data)

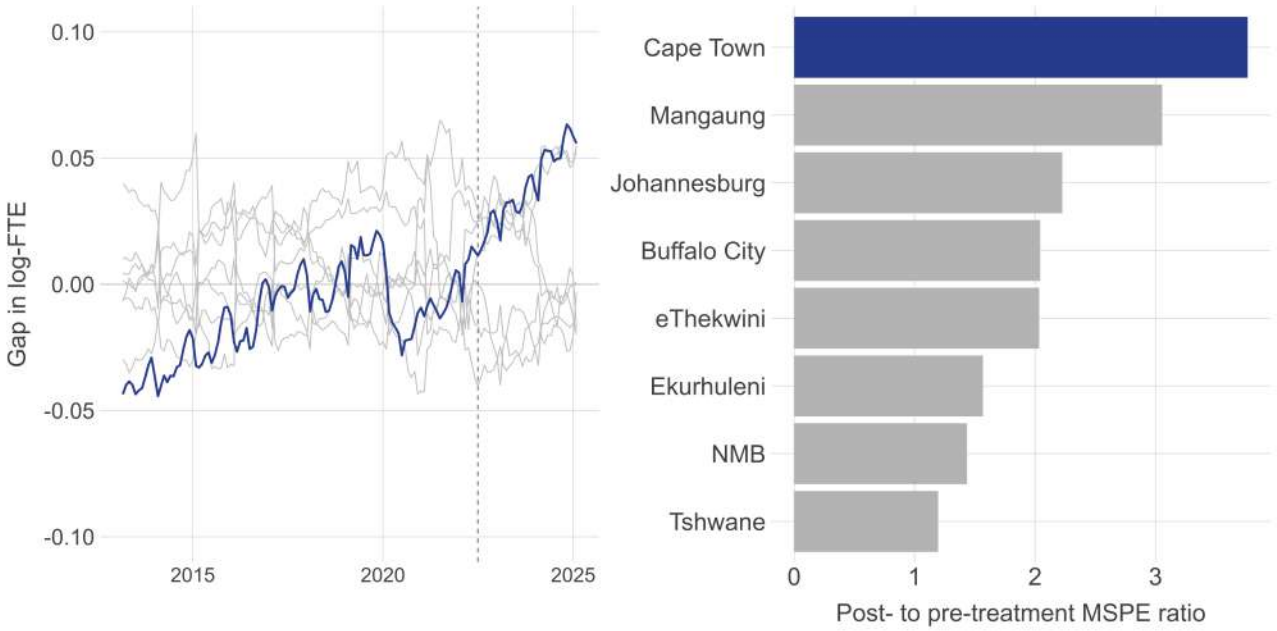

Figure A.3: In-space placebo results (STP-FTE Data)

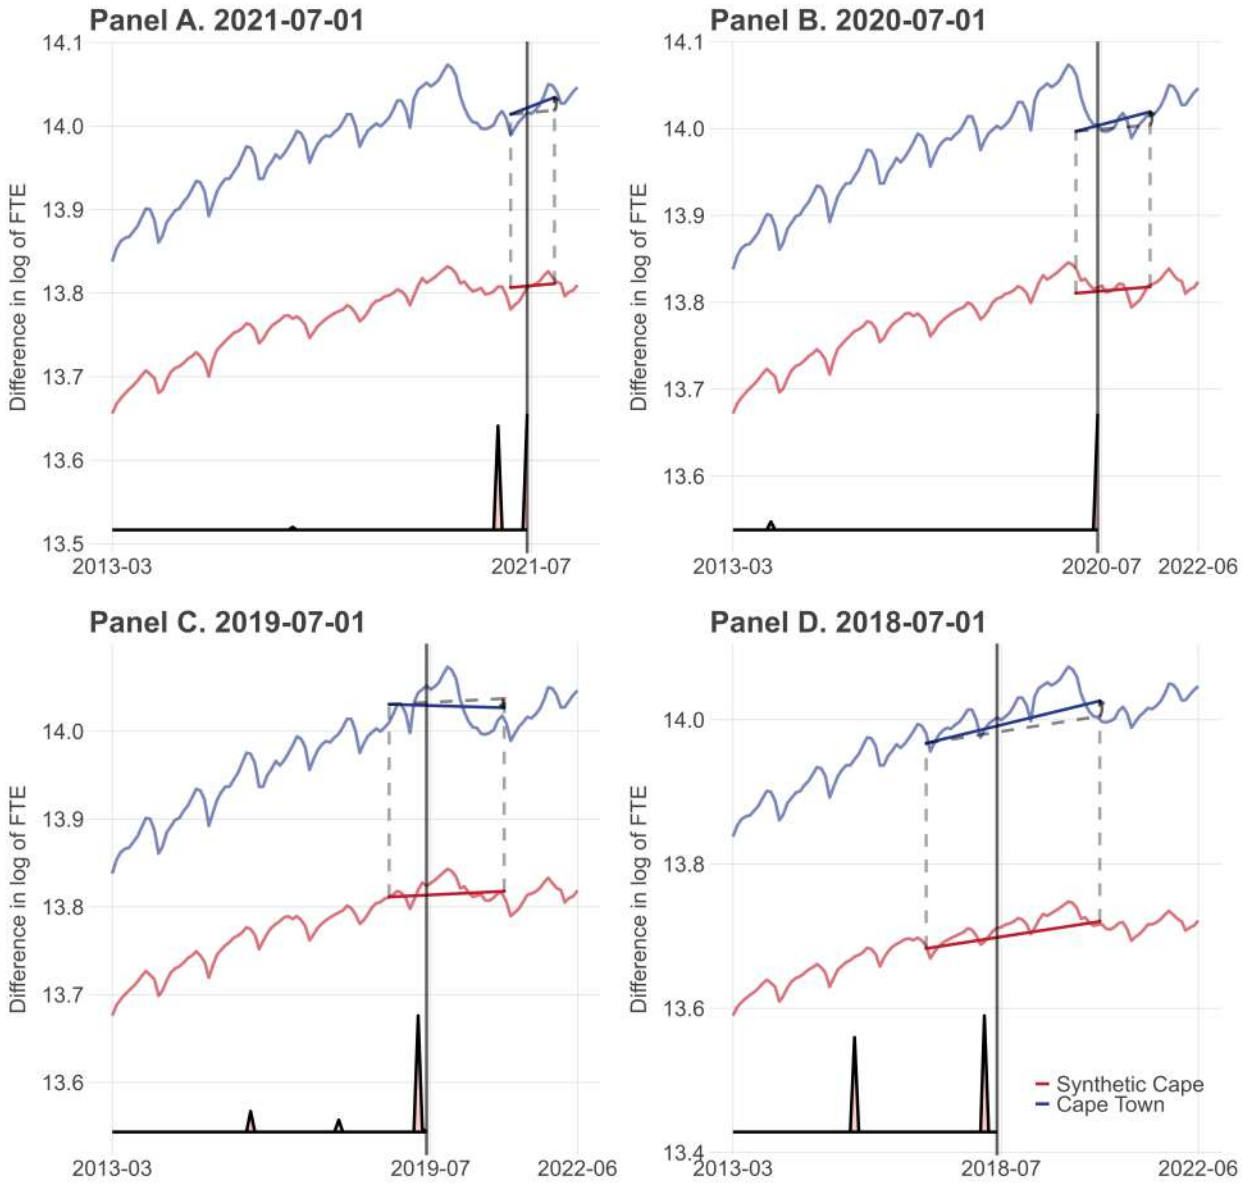

Figure A.4: In-time placebo results (STP-FTE Data)

Table A.1: SIC One-Digit Industry-Level Decomposition

SIC	Industry	Pre FTE	Share	Post FTE	Obs. growth	$\hat{\tau}$	$\widehat{\Delta L}$
1	Agriculture, Forestry and Fishing	23,466	0.019	21,845	−6.9%	−0.010	−228
2	Mining and Quarrying	3,123	0.003	3,203	2.6%	0.052	166
3	Manufacturing	151,265	0.122	148,836	−1.6%	−0.011	−1,649
4	Electricity, Gas and Air Conditioning	7,797	0.006	8,803	12.9%	0.099	810
5	Water Supply, Sewerage and Waste	5,753	0.005	4,234	−26.4%	−0.037	−211
6	Construction	49,737	0.040	54,253	9.1%	0.072	3,733
7	Wholesale and Retail Trade	200,109	0.162	207,227	3.6%	−0.008	−1,544
8	Transportation and Storage	47,531	0.038	46,868	−1.4%	0.012	579
9	Accommodation and Food Services	54,534	0.044	70,012	28.4%	0.044	2,426
10	Information and Communication	51,580	0.042	56,513	9.6%	0.089	4,775
11	Financial and Insurance Activities	82,534	0.067	89,849	8.9%	0.078	6,684
12	Real Estate Activities	14,531	0.012	15,061	3.6%	0.051	756
13	Professional, Scientific and Technical	55,735	0.045	55,706	−0.1%	0.008	459
14	Administrative and Support Activities	131,527	0.106	145,263	10.4%	0.049	6,572
15	Public Administration and Defence	158,762	0.129	157,738	−0.6%	0.018	2,804
16	Education	66,208	0.054	67,415	1.8%	0.012	775
17	Human Health and Social Work	59,824	0.048	60,811	1.7%	0.070	4,313
18	Arts, Entertainment and Recreation	15,554	0.013	18,698	20.2%	0.087	1,415
19	Other Service Activities	50,939	0.041	63,522	24.7%	0.083	4,419
20	Activities of Households as Employers	979	0.001	1,014	3.6%	−0.067	−63
21	Extraterritorial Organizations	3,543	0.003	4,627	30.6%	−0.014	−50
	Total	1,235,031	1.000	1,301,496	—	—	36,940

Disclaimer

This report has been prepared by Codera Analytics (Pty) Ltd ('Codera'). This policy paper is for informational purposes only and does not constitute professional advice. The views expressed are those of the authors and do not necessarily reflect the position of Codera Analytics.

No liability will be accepted by Codera for any loss, damage or expense which may arise directly or indirectly from use of or reliance on any information provided by us or through any act or omission taken as a result of reading or interpreting any of the information provided by us.

If you are not a subscriber, please contact info@codera.co.za to explore the advantages that a subscription can offer.