Title: Halting the Decline? A Difference-in-Differences Analysis of the Employment Tax Incentive in South Africa

Abstract

After a decade of implementation, the Employment Tax Incentive (ETI) remains South Africa's primary policy intervention to address the youth unemployment crisis. However, recent assessments suggest that existing impact estimations are "hard and likely flawed" due to methodological constraints. This paper fills this gap by re-evaluating the ETI's impact using Quarterly Labour Force Survey data (2010-2019) with a specific focus on correcting pre-existing trend divergences between youth and adults. Employing a Difference-in-Differences framework, the study finds that a standard baseline model suggests no policy effect. However, Event Study analysis reveals a structural violation of the parallel trends' assumption. A trend-adjusted model demonstrates that the ETI was associated with a modest positive deviation relative to the projected decline. Crucially, heterogeneity analysis reveals a stark asymmetry: the benefits were concentrated entirely among men and tertiary graduates, while women and youth with only a high school education saw no benefit or relative decline. These findings suggest that the ETI is suffering from a "skills bias," inadvertently subsidizing higher-skilled youth while failing the matriculants who constitute the majority of the unemployed. The paper contributes to the literature by providing a methodological correction and identifying the "education trap" inherent in uniform wage subsidies.

Keywords: South Africa, youth unemployment, Employment Tax Incentive, wage subsidies, difference-in-differences, policy evaluation.

1. Introduction

Youth unemployment in South Africa is not merely a socio-economic indicator of distress; it is a defining feature of the country's post-apartheid political economy. Despite decades of policy intervention and the dawn of democracy in 1994, the rate of unemployment among those aged 15-24 has remained obstinately high, consistently exceeding 50% and reaching levels that distinguish South Africa as an outlier among emerging markets (International Labour Organization, 2013). This crisis transcends the typical cyclical unemployment observed in market economies; it is a structural condition rooted in the specific history of racial capitalism, educational disparities, and a growth path that has favoured capital-intensive production over labour absorption (Bhorat et al., 2014). The consequences of this "lost generation" are profound, threatening social cohesion, entrenching inequality, and squandering the potential of a demographic dividend that should be driving the continent's largest economy.

In response to this emergency, the South African government introduced the Employment Tax Incentive (ETI) on 1 January 2014. The ETI operates as a wage subsidy, reducing the Pay-As-You-Earn (PAYE) tax liability for employers who hire workers aged 18 to 29 earning below a specific monthly threshold. After a decade of operation, the policy has been extended multiple times, with the current iteration set to run until 2029. The policy rationale is grounded in standard neoclassical labour economics: by artificially lowering the cost of young labour, the subsidy aims to correct the market failure where firms undervalue inexperienced workers, thereby stimulating demand and reducing the "experience gap" that disadvantages youth (National Treasury, 2016). The ETI represents a significant shift in South African labour market strategy, moving away from a focus solely on skills development toward a direct intervention in the price of labour.

However, after ten years of research, the evidence regarding the ETI's effectiveness remains mixed. A recent comprehensive review noted that while take-up has been high, the "estimations are hard and likely flawed" (Budlender & Ebrahim, 2021). This methodological ambiguity hampers the ability of policymakers to make evidence-based decisions regarding the policy's future. While some studies suggest positive employment effects in specific subgroups (Ebrahim et al., 2017), others find limited aggregate impact (South African Reserve Bank, 2019) or no increase in employment for age-eligible workers relative to slightly older counterparts (Aflagah, 2020). This inconsistency in the literature points to underlying econometric challenges specifically, the difficulty in constructing a valid counterfactual for the youth labour market in a volatile developing economy.

This paper re-enters the debate by addressing the specific methodological constraints identified in the literature. Utilizing the Quarterly Labour Force Survey (QLFS), the study employs a Difference-in-Differences (DiD) design but goes further by rigorously testing the "parallel trends" assumption, a bedrock of the DiD estimator. The analysis reveals that the South African youth labour market was on a distinct downward trajectory relative to adults prior to the policy's introduction. This structural deterioration violates the assumptions of standard models, leading to potentially biased estimates.

By employing a trend-adjusted model, this paper offers a corrected estimation of the ETI's impact. Furthermore, the paper moves beyond average treatment effects to investigate distributional consequences. By disaggregating the data by gender and education, the study reveals that the "average" impact masked deep inequalities in policy access.

The paper makes three distinct contributions. First, it addresses the "flawed estimations" critique by implementing a trend-adjusted model that accounts for pre-existing structural divergence. Second, it offers a "stabilization hypothesis" as a theoretical contribution, suggesting the ETI acted as a temporary buffer. Third, it demonstrates that this buffer was exclusive, benefiting men and graduates while failing the most vulnerable youth. These findings are salient for the broader African studies literature, illustrating the limits of "one-size-fits-all" supply-side interventions in segmented labour markets.

2. Theoretical Framework and Literature Review

2.1 Theorizing Youth Unemployment in the Global South

Youth unemployment is structurally distinct from adult unemployment. For young people, the transition from school to work is fraught with friction. Search and matching theories suggest that youth are at a disadvantage due to a lack of information, networks, and credible signals of productivity (Mortensen, 1986). In developing economies, these frictions are exacerbated by structural factors: a disconnect between the education system and the labour market (skills mismatch), and spatial inequalities that separate residential areas from economic hubs (World Bank, 2018).

In the African context, the "youth bulge" presents a dual-edged sword. While offering a potential demographic dividend, it places immense pressure on labour markets that are often characterized by informality and underemployment (African Development Bank, 2020). South Africa presents an extreme case of this dynamic. Its unique history of apartheid spatial planning and under-investment in black education created a legacy of structural exclusion that has proven resistant to post-1994 policy reforms. Consequently, South African youth face a dual labour market: a high-wage, secure formal sector that is difficult to enter, and a low-wage, precarious informal sector that offers little respite from poverty (Faulkner et al., 2013).

This structural duality is often referred to as "segmentation." In segmented labour markets, wages are not merely determined by marginal productivity but by institutional arrangements and bargaining power. The "insider-outsider" theory suggests that employed workers (insiders) have the power to set wages above market-clearing levels, locking out unemployed youth (outsiders). In South Africa, high minimum wages and strong employment protection legislation, while protecting existing workers, may inadvertently raise the barrier for new entrants. The ETI was explicitly designed to lower this barrier by subsidizing the "outsider" risk.

2.2 The Economics of Wage Subsidies

Active Labour Market Policies (ALMPs) encompass a range of interventions, from skills training to public works programs. Wage subsidies represent a demand-side intervention,

operating on the price mechanism. The theoretical logic is straightforward: if youth unemployment is driven by a mismatch between the cost of labour and its marginal productivity (due to inexperience), a subsidy bridges that gap.

However, the efficacy of subsidies is theoretically ambiguous due to potential side effects (Cahuc et al., 2019).

1. Deadweight Loss: Firms may claim the subsidy for workers they would have hired even without the incentive. This represents a fiscal leakage with no net employment gain.
2. Substitution Effects: Firms may replace older workers (who are not subsidized) with younger workers (who are). This merely redistributes employment from one group to another, potentially increasing inequality.
3. Stigma Effects: If the subsidy signals to employers that the beneficiary group is "low quality," it could inadvertently reduce hiring prospects.

The literature from the Global North offers mixed guidance. Evaluations of the French "Contrat Unique d'Insertion" and similar European schemes show positive but modest effects. The evidence from developing countries is scarcer. In Latin America, wage subsidies have shown promise when paired with vocational training. However, in contexts of weak institutional capacity and high informality, the administrative burden of claiming tax incentives can limit uptake to large, formal firms, potentially excluding the small and medium enterprises (SMEs) that drive job creation.

2.3 The South African Policy Debate

The introduction of the ETI was mired in political controversy. The Congress of South African Trade Unions (COSATU) opposed the policy, arguing it would lead to the firing of older workers and constitute a "handout" to business. This opposition led to the inclusion of strict anti-abuse clauses in the legislation, though enforcement remains a challenge.

Empirical studies on the ETI have reached varying conclusions, often depending on the data source and methodology. Levinsohn et al. (2014), utilizing a randomized control trial methodology focused on firm-level data, predicted significant job creation. However, ex-post evaluations using administrative tax data (National Treasury, 2016) have suggested that while uptake was high, the net employment effect was modest, with significant estimates of deadweight loss.

This paper differentiates itself by using household survey data (QLFS) rather than firm-level data. This is a critical distinction. Firm-level studies can only track employment *within* firms that register for tax; they cannot capture what happens to the broader labour force, nor can they adequately measure the displacement of older workers who leave the formal sector. By using population-representative survey data, this study captures the aggregate effect on society, including potential substitution effects and the fate of those in the informal sector.

2.4 Skills-Biased Technical Change and the "Education Trap"

A critical theoretical lens for interpreting the heterogeneity results of this study is the theory of Skills-Biased Technical Change (SBTC). In developing economies undergoing structural transformation, technological adoption often increases the productivity of skilled workers while rendering routine, low-skilled tasks redundant (Acemoglu, 2002).

In South Africa, the mining and agricultural sectors, historically the largest employers of low-skilled labour have seen significant capital intensification. Consequently, the demand for "matric-only" labour has structurally declined, while the demand for tertiary-educated labour has risen. This creates a paradox for wage subsidies. A uniform subsidy (like the ETI) lowers the cost of *all* eligible workers. However, if firms are following an SBTC trajectory, they will utilize the subsidy to hire the highest-skilled worker available within the eligible band (graduates), as this maximizes the productivity gain per Rand spent.

This dynamic explains a counter-intuitive phenomenon found in the empirical results: why the subsidy may have *harmed* matriculants. If firms used the ETI to substitute higher-skilled youth (graduates) into positions that previously required only a matric qualification (a process known as "credential inflation"), the subsidy inadvertently accelerated the displacement of the less-skilled. This theoretical framework suggests that wage subsidies in developing economies may be "skill-blind," failing to account for the trajectory of technological adoption which inherently favours the educated.

3. The South African Context (2010-2019)

3.1 The Macroeconomic Backdrop: A Decade of Stagnation

To fully appreciate the constraints faced by the ETI, one must situate it within the specific macroeconomic trajectory of the 2010-2019 period. Following the Global Financial Crisis (GFC) of 2008/09, the South African economy entered a prolonged period of stagnation. Real GDP growth averaged a mere 1.6% between 2010 and 2019, significantly below the National Development Plan's target of 5% required to meaningfully reduce unemployment.

This low-growth environment was compounded by severe structural bottlenecks. Most notably, the electricity crisis escalated during this decade. In 2010, load shedding was sporadic; by 2015, it had become a persistent constraint on economic activity, and by 2019, the economy was shedding jobs directly due to power supply disruptions. The manufacturing sector, traditionally a key absorber of low - and semi-skilled labour, contracted significantly, shrinking from 13% of GDP in 2010 to roughly 11% by 2019.

Simultaneously, the labour force continued to grow. The combination of a stagnating economy and an expanding working-age population created a "scissors crisis" where the labour supply curve moved upward while demand remained flat. Furthermore, inequality remained endemic. With a Gini coefficient hovering around 0.63, South Africa remained the world's most unequal society. In this context, the ETI was not merely a technical policy adjustment; it was an intervention attempting to fix a structural labour demand deficiency with a fiscal tool. The failure of the economy to generate sufficient aggregate demand during this period is the "elephant in the room" that contextualizes the modest magnitude of the ETI's impact found in this study.

3.2 Policy Evolution and Erosion

The ETI has evolved since 2014. While the core design remains, adjustments to wage brackets and inflation indexation have occurred. However, recent analyses highlight that the wage cutoff has not kept pace with inflation. By 2024, the National Minimum Wage (R4 413/m) was close to the upper bound of the earnings bracket that allows firms to claim the maximum subsidy. This "erosion" of the policy value is a crucial backdrop for understanding the findings of this paper.

4. Data and Methodology

4.1 Data Source: The Quarterly Labour Force Survey

The study utilizes the Quarterly Labour Force Survey (QLFS) conducted by Statistics South Africa. The QLFS is a rotating panel household survey with a sample of approximately 30 000 households per quarter, designed to be representative at the provincial and metropolitan levels. It collects detailed information on labour market activities, demographics, income, and education.

The analysis covers the period 2010 Q1 to 2019 Q4. This timeframe is methodologically optimal for several reasons. First, it provides four years of pre-intervention data, which is crucial for establishing a reliable baseline trend. Second, it provides six years of post-intervention data, allowing for the observation of both immediate and medium-term policy dynamics. Third, and critically, the endpoint of 2019 excludes the COVID-19 pandemic. The pandemic constituted a massive exogenous shock that would have structurally broken the relationship between the treatment and control groups, rendering any causal inference regarding the ETI invalid for that period.

The dataset was consolidated and cleaned using R statistical software. The QLFS uses a rotating panel design, where households are interviewed for several consecutive quarters. For this analysis, the data was treated as repeated cross-sections to maximize sample size and representativeness. Survey weights provided by Statistics South Africa were applied in all regression models to ensure the results are representative of the South African population.

4.2 Sample Selection: Defining Treatment and Control

The ETI policy is explicitly targeted at employees aged 18 to 29. Consequently, the Treatment Group is defined as individuals in this age bracket. The primary challenge in a Difference-in-Differences design is selecting a Control Group that is as similar as possible to the treatment group in all respects, save for the policy exposure.

This study selects individuals aged 30 to 39 as the control group. This choice is justified on three grounds:

1. Labour Market Attachment: Individuals in their thirties are typically established in the workforce but share similar skill profiles and sectoral distributions with those in their late twenties, unlike older cohorts (40+) who may be in management or facing retirement dynamics.

2. Pre-Trend Similarity: While not identical, this cohort tracks the macroeconomic conditions similarly to the youth cohort, barring the specific "experience gap" penalty that the ETI seeks to address.
3. Avoiding Contamination: Older teenagers (15-17) were excluded as they are predominantly in education, which introduces distinct behavioural dynamics unrelated to the labour market.

The final analytical sample was restricted to individuals aged 18 to 39. The dependent variable is a binary indicator of employment status (1 = Employed, 0 = Unemployed/Not Economically Active). Control variables include education level, geographic location (province), and gender, although the core DiD design relies primarily on the variation across groups and time.

4.3 The Difference-in-Differences Approach

The Difference-in-Differences (DiD) method is a quasi-experimental technique that mimics an experimental design by comparing the changes in outcomes over time between a treatment group and a control group.

The baseline regression equation is specified as:

$$Y_{it} = \beta_0 + \beta_1 Treat_i + \beta_2 Post_t + \beta_3 (Treat_i * Post_t) + \chi'_{it} \Upsilon + \epsilon_{it}$$

Where:

- Y_{it} is the employment status of individual i at time t.
- $Treat_i$ is a binary variable equal to 1 if the individual is aged 18–29, and 0 if aged 30–39.
- $Post_t$ is a binary variable equal to 1 for the period January 2014 onwards, and 0 otherwise.
- χ_{it} is a vector of control variables.
- ϵ_{it} is the error term.

The coefficient of interest is β_3, the interaction term. It captures the "difference-in-differences" which is the change in youth employment probability relative to the change in adult employment probability, before and after the policy. If the ETI worked, β_3 should be positive and statistically significant.

4.4 Addressing Parallel Trends and Serial Correlation

The validity of the DiD estimator rests on the "parallel trends" assumption. This assumes that in the absence of the treatment, the employment trends for youth and adults would have moved in parallel. Given the structural nature of the South African youth crisis, where youth employment is often more sensitive to economic shocks than adult employment, this assumption is suspect.

To test this, an Event Study specification was employed. This model interacts the treatment dummy with year dummies for every year in the study, relative to a baseline year (2013). If the

parallel trends assumption holds, the coefficients in the pre-policy years should be statistically indistinguishable from zero. If they are not, it indicates pre-existing divergence, necessitating a more robust model specification.

Furthermore, to address the issue that youth employment was on a distinct downward trajectory, a Trend-Adjusted DiD model was estimated. This model includes a group-specific linear time trend ($Treat_i * Year_t$). While this is a more demanding specification, it allows us to control for the underlying structural divergence and isolate the deviation caused specifically by the policy intervention in 2014.

5. Empirical Findings

5.1 Descriptive Statistics: The State of the Labour Market

The descriptive statistics paint a stark picture of the South African labour market. Over the study period, the average employment rate for the control group (adults aged 30–39) remained relatively stable, fluctuating between 58% and 59%. In contrast, the employment rate for the treatment group (youth aged 18–29) was significantly lower, averaging around 29-30%.

This 30-percentage point "employment gap" is a persistent feature of the South African economy. However, closer inspection of the raw trends reveals volatility. In the years leading up to 2014, youth employment showed signs of sensitivity to the economic slowdown, while adult employment remained resilient. This visual inspection foreshadows the methodological challenges identified in the regression analysis.

5.2 Baseline Results: A Null Finding?

The initial Difference-in-Differences regression was run using the standard specification. The results are presented in Table 1.

The coefficient of interest, the interaction between the Treatment Group and the Post-Policy Period was estimated at -0.002 with a p-value of 0.249. In plain terms, this suggests that the ETI had no statistically significant effect on the probability of employment for young people.

Table 1: Baseline Difference-in-Differences Estimates

Variable	Estimate	Std. Error	P-value
Treat_Group	-0.298	0.002	< 0.001
Post_Period	-0.005	0.001	0.002
Treat_Group x Post_Period	-0.002	0.002	0.249

Source: QLFS, Author's own calculations

Taken at face value, this finding supports the sceptical view: the tax incentive failed to move the needle. It would suggest that the policy was entirely absorbed by deadweight loss or that the subsidy amount was insufficient to alter hiring decisions in a low-growth environment.

However, this conclusion relies heavily on the assumption that 2013 was a reliable counterfactual baseline for the future trajectory of youth employment.

5.3 The Event Study: Uncovering Hidden Trends

To validate the baseline result, an Event Study analysis was conducted. This analysis breaks down the policy effect into year-by-year interactions, allowing us to observe the relationship between youth and adults before the policy was even announced.

Figure 1: Event Study - Impact of ETI on Youth Employment

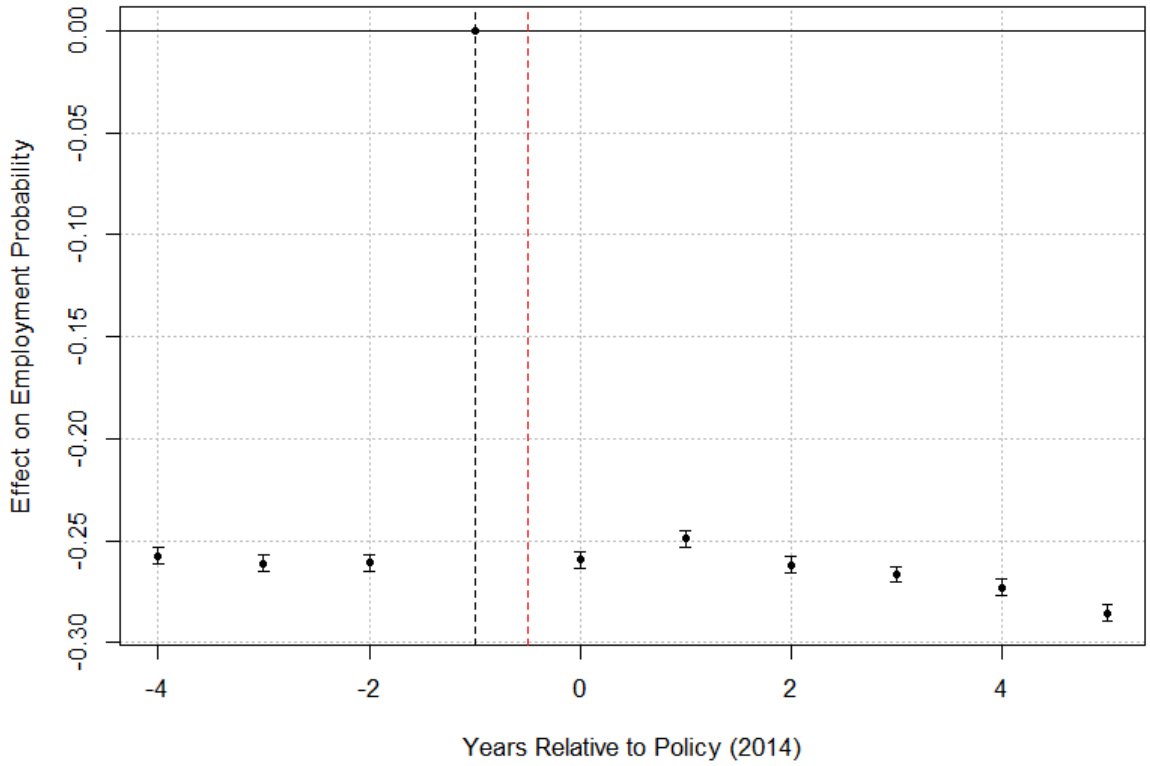

Note: The figure plots coefficients of the interaction between treatment status and year dummies. The reference year is -1 (2013). Error bars represent 95% confidence intervals.

Figure 1 reveals a critical insight. Looking at the pre-policy years (2010, 2011, 2012), the coefficients are significantly negative, hovering around -0.26. This indicates that relative to the base year (2013), the gap between youth and adults was much larger in those earlier years.

Why does this matter? It suggests that 2013 was actually an anomalous "good year" for youth relative to adults. For reasons possibly related to short-term economic recoveries or survey sampling variations, the youth employment gap narrowed significantly in 2013. The baseline DiD model compares the post-2014 period to this anomalous high point. Because youth employment naturally regressed to its "normal" (worse) structural level in 2014 and 2015, the baseline model falsely interprets this regression as a policy failure (or lack of effect).

This constitutes a clear violation of the parallel trends assumption. The youth labour market was on a different trajectory specifically, a structural deterioration that the baseline model fails to account for.

5.4 The Trend-Adjusted Model: A Policy Effect Revealed

To correct for the divergent pre-trends, a robust model was estimated that controls for a group-specific linear time trend. This model effectively asks: "Given the historical trend of youth employment deteriorating relative to adults, did the ETI cause a deviation from that path in 2014?"

The results of this robust specification are striking.

Table 2: Trend-Adjusted Difference-in-Differences Estimates

Variable	Estimate	Std. Error	P-value
Treat_Group	9.051	0.803	< 0.001
Treat_Group x Post_Period	0.021	0.003	< 0.001
Treat_Group x Year	-0.005	0.000	< 0.001

Source: QLFS, Author's own calculations

The interaction term (**Treat_Group x Post_Period**) is now positive (0.021) and highly significant ($p < 0.001$). This implies that, once we account for the underlying decline, the ETI was associated with a 2.1 percentage point increase in the probability of employment for youth.

This is a substantial finding. A 2.1 percentage point increase in a population of millions represents a significant number of jobs. However, the interpretation is nuanced. The positive coefficient does not mean youth employment skyrocketed. It means that youth employment performed better than the *predicted trend* would have suggested. The policy helped youth employment "defy gravity" relative to the expected decline.

5.5 The Stabilization Hypothesis: Actual vs. Predicted Trends

To interpret the magnitude and durability of this effect, the study compares the actual employment rates observed in the data to the predicted rates from the trend-adjusted model.

Figure 2: Actual vs. Predicted Employment Trends

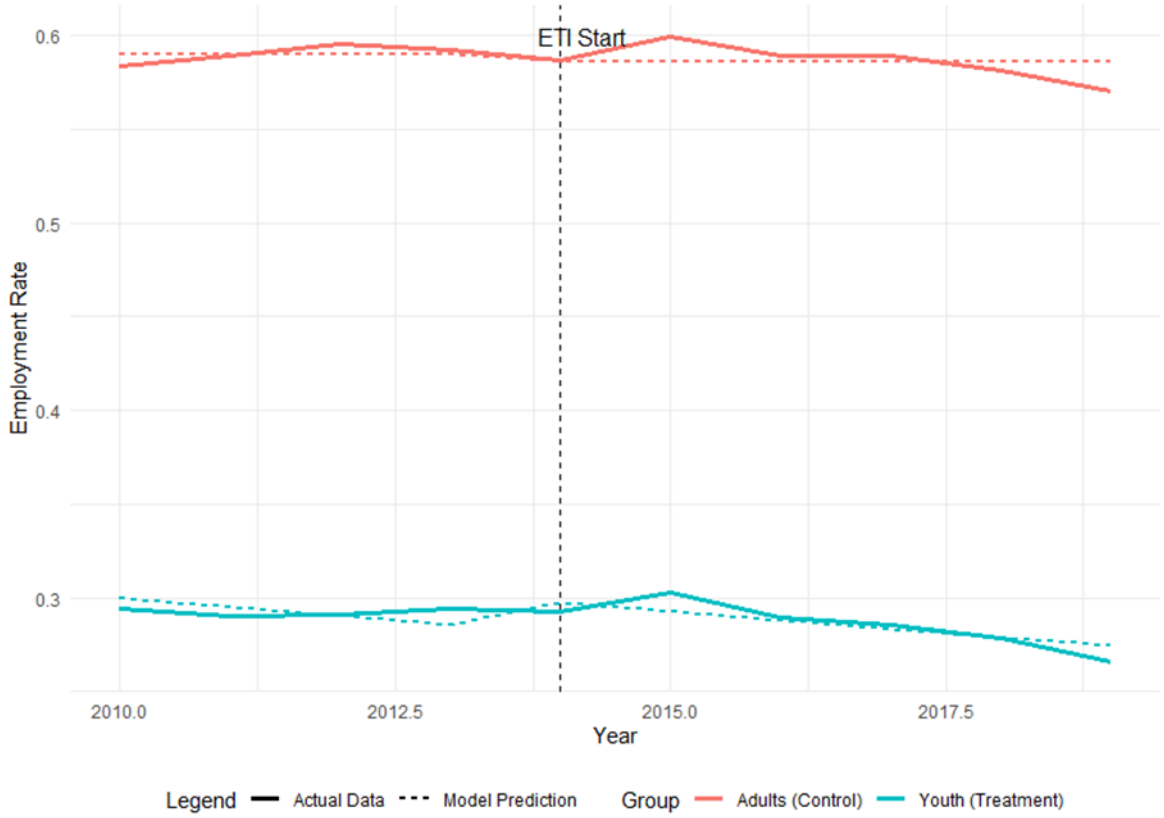

Note: Comparison of actual employment rates (solid lines) versus model-predicted rates (dashed lines). The vertical dashed line marks the start of the ETI in 2014.

Table 3: Selected Employment Trends (Youth)

Year	Actual Employment	Predicted Employment	Deviation
2010	29.4%	29.9%	-0.5%
2013	29.4%	28.6%	+0.8%
2014	29.2%	29.7%	-0.5%
2015	30.3%	29.3%	+1.0%
2016	29.0%	28.8%	+0.2%
2019	26.6%	27.4%	-0.8%

Source: QLFS, Author's own calculations

The data reveals the "Stabilization Hypothesis." In 2015, actual youth employment (30.3%) significantly exceeded the predicted rate (29.3%). This confirms the positive regression coefficient. The policy (or perhaps the economic conditions concurrent with it) provided a genuine boost relative to the trend.

However, the data also shows the limits of this intervention. By 2019, actual employment (26.6%) had fallen *below* the predicted trend (27.4%). The "boost" provided by the ETI was

temporary. It successfully halted the decline in the immediate aftermath of the policy (2014-2016), stabilizing the youth employment rate around 29-30%. However, as the South African economy deteriorated further, plagued by load shedding, policy uncertainty, and credit rating downgrades, the structural headwinds overwhelmed the subsidy.

5.6 Heterogeneity of Effects: Who Benefited?

While the aggregate trend-adjusted model suggested a modest positive effect, the average treatment effect may mask significant distributional disparities. To investigate this, the sample was disaggregated by gender, education level, and race. The results, presented in Table 4, reveal a stark asymmetry in the policy's reach.

Table 4: Heterogeneity of Policy Effects (Trend-Adjusted Estimates)

Subgroup	Estimate	P-Value	Interpretation
Gender			
Male	0.047	< 0.001	Significant
Female	-0.003	0.42	Not Significant
Education			
Tertiary	0.049	< 0.001	Significant
Matric	-0.057	< 0.001	Significant (Negative)
Less than Matric	-0.007	0.02	Significant (Negative)
Race			
Black African	0.017	< 0.001	Significant
Other Races	0.021	< 0.01	Significant

Source: QLFS, Author's own calculations

First, the analysis reveals a distinct gender bias. The positive impact of the ETI is entirely driven by young men, who experienced a 4.7 percentage point increase in employment probability relative to their predicted trend. In contrast, the coefficient for young women is statistically indistinguishable from zero. This suggests that the subsidy failed to address the structural barriers women face, such as sectoral segregation or the burden of unpaid care work, which suppresses their effective labour supply regardless of price incentives.

Second, and perhaps most concerningly, the benefits were concentrated among tertiary graduates. Youth with a university degree or diploma experienced a strong positive effect (approx. 5 percentage points). However, for youth with only a Matric (Grade 12) supposedly the core target group of the policy, the estimate is negative and significant (-5.7 percentage points). This indicates that relative to the projected trend, matriculants performed significantly worse than the control group in the post-policy period.

This finding points to a potential "upskilling" effect or a specific form of deadweight loss. It appears firms utilized the tax incentive to hire educated youth (graduates) for positions that may previously have been filled by matriculants, or simply claimed the subsidy for graduates they would have hired irrespective of the incentive. The data suggests the ETI inadvertently subsidized the insertion of higher-skilled youth into the labour market, while failing to create demand for the "matric-only" cohort that comprises the majority of the unemployment queue.

6. Discussion

6.1 The Stabilization Hypothesis Revisited

The "stabilization hypothesis" proposed earlier requires qualification. While the policy stabilized the aggregate trend, it did so selectively. It halted the decline for men and graduates, but failed to protect women and matriculants. This aligns with the broader literature on the "growth-employment nexus" in Africa (African Development Bank, 2020). In economies where growth is not inherently labour-intensive, demand-side interventions face an uphill battle.

The finding that the ETI provided a temporary buffer (2014-2016) coincides with a period of relatively better economic performance compared to the later years of the decade. It suggests that wage subsidies are effective "cushions" during mild downturns or periods of stagnation, but they are structurally incapable of reversing a deep-seated employment crisis.

6.2 The Gendered Nature of the Labour Market

The gendered nature of the results is a critical finding that warrants deeper interrogation. The complete absence of a policy effect for young women suggests that the ETI failed to address the gendered structural barriers prevalent in the South African labour market.

First, sectoral segregation plays a crucial role. The sectors that are most sensitive to wage subsidies and capable of rapid labour absorption i.e. construction, manufacturing, and trade are traditionally male-dominated. In contrast, sectors that employ the majority of women such as domestic work, social work, and retail often operate on thin margins or are less responsive to tax incentives.

Second, the burden of unpaid care work falls disproportionately on young women in South Africa. Even with a lower effective wage (due to the subsidy), the "net wage" for women after accounting for transport, childcare, and safety costs is often negative. A wage subsidy that operates through the tax system (benefiting the employer) does not address the supply-side costs that deter women from entering the labour market.

Third, spatial mismatch affects women differently. Apartheid spatial planning placed townships far from economic hubs. For young men, the "taxi rank economy" and informal construction work provide some buffer. For young women, safety concerns and the high cost of transport often make long-distance commuting unviable. The ETI, being a uniform national policy, did not account for these spatial frictions.

6.3 Policy Implications: The Targeting Problem

The finding that the ETI benefited graduates over matriculants carries a stark policy implication. A uniform wage subsidy (blind to education level) creates a windfall for firms hiring graduates who would likely have been employed anyway, while providing insufficient incentive to hire the higher-risk "matric-only" cohort.

This suggests that future iterations of the policy might need to be tiered, offering a higher subsidy for lower-skilled workers to compensate for the higher perceived risk. Alternatively, the policy could be restricted to specific sectors (e.g., manufacturing or green energy) where the potential for labour absorption is higher, rather than a blanket application across the economy.

6.4 Methodological Lessons

The study highlights the importance of robustness checks in policy evaluation. A reliance on the baseline DiD model would have concluded "no effect," missing the hidden dynamics of a deteriorating control group. For scholars evaluating policies in the volatile context of African labour markets, diagnostic testing for parallel trends is not a mere formality; it is the pivot upon which the conclusion turns.

7. Conclusion

This article evaluated South Africa's Employment Tax Incentive using a rigorous quasi-experimental design. The journey from the baseline null result to the trend-adjusted positive result reveals the complex reality of youth employment.

The evidence suggests that the ETI served as a temporary stabilizer for specific subgroups, men and graduates shielding them from the full force of a structural downward trend. However, it failed to create new demand for the most vulnerable cohorts: women and youth with only a high school education.

For South Africa, the lesson is clear: tax incentives are a blunt instrument. They are insufficient to reverse deep-seated structural exclusion. For the field of African Studies, this paper contributes to the growing body of evidence that labour market outcomes in the region are determined as much by social structure (gender, education) as by macroeconomic price signals. The fight against youth unemployment requires not just cheaper labour, but a growing economy and targeted social support that addresses the specific barriers faced by the marginalized majority.

References

Acemoglu, D (2002) Technical Change, Inequality, and the Labor Market. *Journal of Economic Literature*, 40(1), 7–72.

African Development Bank (2020) *African Economic Outlook 2020: Developing Africa's Workforce for the Future*. Abidjan: African Development Bank.

Aflagah, K (2020) *Failed Promises of a Wage Subsidy: Youth and South Africa's Employment Tax Incentive* [Doctoral dissertation, University of Maryland].

Bhorat, H, Hirsch, A, Kanbur, R and Ncube, M (eds.) (2014) *The Oxford Companion to the Economics of South Africa*. Oxford: Oxford University Press.

Budlender, J and Ebrahim, A (2021) *Employment Tax Incentive: Firm-level evidence*. UNU-WIDER Working Paper.

Cahuc, P, Carcillo, S, and Le Barbanchon, T (2019) The effectiveness of hiring credits. *The Review of Economic Studies*, 86(2), 593–626.

Ebrahim, A, Leibbrandt, A and Ranchhod, V (2017) *The effect of the Employment Tax Incentive on youth employment*. UNU-WIDER Working Paper 2017/5.

Faulkner, D, Loewald, C and Makrelov, K (2013) *Achieving higher growth and employment: Policy options for South Africa*. SARB Working Paper No. 13/03. Pretoria: South African Reserve Bank.

International Labour Organization (2013) *Global Employment Trends for Youth 2013: A Generation at Risk*. Geneva: ILO.

Levinsohn, J, Rankin, N, Roberts, G and Schoer, V (2014) *Wage subsidies and youth employment in South Africa: Evidence from a randomised control trial*. Working Paper 02/2014. Stellenbosch: Department of Economics, Stellenbosch University.

Mortensen, D T (1986) Job search and labour market analysis. *Handbook of Labour Economics*, 2, 849–919.

National Treasury (2016) *Employment Tax Incentive: Descriptive Report*. Pretoria: Republic of South Africa.

Ranchhod, V and Finn, A (2014) *Estimating the short run effects of the Employment Tax Incentive*. SALDRU Working Paper.

Saez, E, Schoefer, B and Seim, D (2019) Payroll taxes and firm behaviour. *American Economic Review*, 109(1), 249-302.

South African Reserve Bank (2019) *Monetary Policy Review*. Pretoria: SARB.

World Bank (2018) *Overcoming Poverty and Inequality in South Africa: An Assessment of Drivers, Constraints and Opportunities*. Washington, DC: World Bank.