

The Credential Cliff: Education Thresholds and Labour Market Gaps in South Africa

APEP Autonomous Research*

@olafdrw

March 9, 2026

Abstract

In South Africa, a single percentage point on a matric exam can determine whether a young person gains access to university, a diploma programme, or nothing at all. This paper documents the “credential cliff” — large descriptive gaps in labour market outcomes across education thresholds — using aggregate data from the National Senior Certificate (2008–2022), the Quarterly Labour Force Survey, and cross-country comparisons from 19 nations. The absorption rate differs by 20 percentage points between matric completion and a post-school credential, and by another 15 points between a diploma and a university degree. South Africa’s education premium (17 pp) is the largest among comparable middle-income economies. I develop a multi-cutoff regression discontinuity blueprint exploiting the three mechanically assigned pass levels (30%, 40%, 50%) and characterize the institutional conditions for future causal estimation.

JEL Codes: I26, J24, O15

Keywords: credential gaps, regression discontinuity, matric examination, youth unemployment, South Africa

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

Every year, roughly 600,000 young South Africans sit for the National Senior Certificate examination — the “matric” — knowing that their scores will sort them into one of three credential tiers: Higher Certificate, Diploma, or Bachelor’s pass. A student whose fourth-best subject score lands at 49 percent receives a Diploma-eligible credential; one point higher, 50 percent, and she qualifies for university admission. No committee reviews her application. No teacher exercises judgment. The threshold is mechanical, sharp, and life-altering, because South Africa’s labour market treats these credential categories as profoundly different signals of employability.

How different? Consider the raw numbers. Among working-age adults with only a matric, 39 percent are employed. Among those who hold a post-school certificate or diploma, that figure is 59 percent — a 20 percentage-point gap for what is often a two-year programme ([Statistics South Africa, 2019](#)). For university graduates, the absorption rate reaches 74 percent. In a country where one in three young people is unemployed and one in two is not in education, employment, or training, these gaps carry enormous welfare consequences ([Banerjee et al., 2008](#); [Bhorat et al., 2015](#)). A matric pass alone — the credential that the education system was designed around — is associated with remarkably poor labour market outcomes.

This paper contributes to the literatures on education and labour markets, credential signaling, and South African economic policy by (1) documenting large descriptive gaps across credential tiers and (2) developing a complete multi-cutoff RDD blueprint for future causal estimation. The paper makes three specific contributions. First, it documents the institutional architecture of South Africa’s matric pass-level system and shows that it generates a textbook setting for a multi-cutoff regression discontinuity design (RDD). The three pass thresholds — 30 percent for Higher Certificate, 40 percent for Diploma, and 50 percent for Bachelor’s — are mechanically assigned based on subject scores, with no discretion by examiners or moderators ([Department of Basic Education, 2022](#)). The running variable (the binding-constraint subject score at each cutoff) is continuous, and the treatment (credential assignment) is deterministic conditional on the score. These features satisfy the core requirements of the RDD framework ([Lee and Lemieux, 2010](#); [Imbens and Lemieux, 2008](#)).

Second, using publicly available aggregate data — the Department of Basic Education’s NSC Technical Reports (2008–2022), Statistics South Africa’s Quarterly Labour Force Survey, the Department of Higher Education and Training’s enrolment statistics, and the World Bank’s World Development Indicators for 19 comparator countries — I document the gradient of labour market gaps across credential categories. The evidence reveals a “credential cliff”:

outcomes differ modestly within matric pass types (Higher Certificate holders have a 31 percent absorption rate versus 37 percent for Diploma-eligible passers) but differ sharply when individuals acquire a post-school qualification. This 20 percentage-point gap dwarfs the 5.5-point difference between the lowest and middle matric tiers. The pattern is consistent with a model in which employers use the post-school credential as a screening device that matric pass types, by themselves, cannot substitute for (Spence, 1973).

Third, I place South Africa in cross-national perspective. Among 19 middle-income and upper-middle-income economies, South Africa exhibits the largest positive education premium: the gap between the overall unemployment rate and the unemployment rate among tertiary-educated workers is 17 percentage points. In most comparator countries, education confers a modest advantage or none at all (the median premium is approximately zero), while in several — including Nigeria, Tunisia, and Egypt — tertiary-educated workers actually face *higher* unemployment than the general population. South Africa is a striking outlier: the *relative* advantage of education is unusually large, suggesting that credential-based screening plays a distinctively important role in this labour market (Psacharopoulos and Patrinos, 2018).

I am transparent about what this paper cannot do. The ideal implementation of the multi-cutoff RDD requires individual-level examination scores linked to labour market outcomes. Such data exist: DataFirst at the University of Cape Town curates the NSC Examination Database (DOI: 10.25828/pcn8-pc32), which contains individual scores for every candidate from 2008 onward. Linking these to the HEMIS tertiary enrolment records or NIDS panel data would permit estimation of causal effects at each threshold. I lack licensed access to these microdata. Accordingly, I develop the identification framework in full — specifying the running variable, the bandwidth-selection procedure, the multi-cutoff pooling estimator of Cattaneo et al. (2020), and the battery of validity tests (McCrary density tests, covariate balance, donut-hole specifications, cross-cutoff placebos) — but I estimate the credential gradient using aggregate tabulations rather than individual-level regressions. The aggregate evidence is suggestive, not causal. The paper’s value lies in documenting the institutional setting with precision, quantifying the descriptive gradient, and providing a complete blueprint for future individual-level estimation.

The existing literature on education and labour markets in South Africa is substantial but has not exploited the matric pass-level thresholds as a source of quasi-experimental variation. Van der Berg (2007) and Spaull (2015) document the deep quality inequalities inherited from apartheid and show that school quality remains the primary determinant of educational outcomes. Lam et al. (2011) and Branson et al. (2019) trace the school-to-work transition and find that the matric credential, while necessary, is increasingly insufficient for labour market

success. [Ranchhod and Finn \(2021\)](#) estimate returns to matric completion but treat it as a single credential, without distinguishing among pass types. [Kerr et al. \(2019\)](#) provide updated Mincerian returns across education levels but rely on selection-on-observables assumptions.

Internationally, the closest analogues come from settings where centralized examinations sort students into tiered educational tracks. [Pop-Eleches and Urquiola \(2013\)](#) exploit school-assignment cutoffs in Romania to estimate effects on achievement and university attendance. [Jackson \(2010\)](#) uses rule-based assignment in Trinidad and Tobago to study school quality effects. [Hastings et al. \(2013\)](#) and [Kirkeboen et al. \(2016\)](#) estimate returns to specific fields and institutions in Chile and Norway, respectively, using admission cutoffs. [Zimmerman \(2014\)](#) estimates returns to college admission for marginal students in Florida. Each of these studies exploits a single cutoff or a set of institution-specific cutoffs. The South African matric system offers something distinctive: three national cutoffs, applied uniformly to every candidate, that map directly onto credential tiers with clear labour market interpretations.

The broader returns-to-education literature provides a foundation for interpreting the credential gradient. [Card \(1999\)](#) surveys instrumental-variables estimates and finds returns of 8–13 percent per year of schooling. [Duflo \(2001\)](#) exploits school construction in Indonesia and estimates returns of 6–10 percent. [Oreopoulos \(2006\)](#) uses compulsory-schooling reforms in the United Kingdom to estimate returns of 12–16 percent. [Carneiro et al. \(2011\)](#) show that marginal returns differ substantially from average returns, with the largest gains accruing to individuals at the margin of completing additional schooling. This heterogeneity is precisely what a multi-cutoff RDD in South Africa could identify: students near each threshold are the relevant marginal population for each credential tier.

The paper also speaks to the signaling-versus-human-capital debate. In a pure human-capital model ([Becker, 1964](#); [Mincer, 1974](#)), the matric pass type matters only insofar as it determines access to further education, and the returns to crossing each threshold should be proportional to the additional human capital acquired through the unlocked educational pathway. In a signaling equilibrium ([Spence, 1973](#)), the credential itself carries informational value, and crossing a threshold can change labour market outcomes even absent additional education. The three-cutoff structure permits a partial test: if crossing the Bachelor’s threshold improves employment even among those who do not enrol in university, the signaling channel is operative. The aggregate data cannot isolate this channel, but I show that the institutional conditions for such a test are well-defined.

Section 2 describes the institutional setting; Sections 3–5 develop the conceptual and empirical framework; Sections 6–7 present descriptive results and robustness checks; and Sections 8–9 discuss implications and conclude.

2. Institutional Background and Policy Setting

2.1 The National Senior Certificate Examination

The National Senior Certificate (NSC) examination, commonly known as the “matric,” is South Africa’s school-leaving examination. Introduced in its current form in 2008, it replaced the earlier Senior Certificate and is administered by the Department of Basic Education (DBE) to all Grade 12 learners in the public school system. Approximately 600,000 to 700,000 candidates write the examination annually, making it one of the largest standardized assessments in sub-Saharan Africa ([Department of Basic Education, 2022](#)).

The examination is high-stakes in a precise sense: the matric credential is the minimum requirement for formal-sector employment in most industries and the sole gateway to post-secondary education. Unlike systems in which university admission depends on a combination of school grades, teacher recommendations, and entrance exams, South Africa’s post-secondary access is determined almost entirely by the matric pass type and the aggregate percentage score. The examination therefore serves simultaneously as a school-leaving assessment, a university entrance examination, and a labour market credential.

2.2 The Three Pass Levels

The NSC assigns learners to one of three hierarchical pass levels based on their subject scores. The rules are codified in the *National Policy Pertaining to the Programme and Promotion Requirements of the National Curriculum Statement, Grades R–12* and applied mechanically by the Umalusi certification council. The three levels are:

1. **Higher Certificate Pass (30% threshold):** A learner must achieve at least 30 percent in the language of instruction, 30 percent in two other subjects, and pass at least three additional subjects with 20-credit status. This is the minimum matric pass and qualifies the holder for Higher Certificate programmes at TVET colleges.
2. **Diploma Pass (40% threshold):** A learner must achieve at least 40 percent in the language of instruction, at least 40 percent in three subjects of 20-credit status (excluding Life Orientation), and at least 30 percent in two additional subjects. This qualifies the holder for diploma programmes at universities of technology and comprehensive universities.
3. **Bachelor’s Pass (50% threshold):** A learner must achieve at least 50 percent in four subjects of 20-credit status (excluding Life Orientation), at least 30 percent in the language of instruction, and at least 30 percent in two additional subjects. This is the

only credential that qualifies the holder for admission to bachelor’s degree programmes at universities.

The critical feature for identification is that these thresholds are *mechanical*. No examiner, teacher, or moderator exercises discretion over the pass-level assignment. A learner whose fourth-best 20-credit subject score is 49 percent cannot receive a Bachelor’s pass; one whose score is 50 percent automatically does. The running variable — the subject score that constitutes the binding constraint at each cutoff — is continuous and finely measured (marks are recorded as whole percentages, yielding a discrete grid with approximately 70 mass points in the relevant range).

2.3 Score Assignment and Moderation

The examination is set nationally by the DBE, with subject-specific examination panels developing question papers. Marking follows a standardized memorandum, with external moderators reviewing a sample of scripts. Umalusi, the quality council for general and further education, conducts a statistical moderation process after marking is complete. This moderation can adjust raw marks upward or downward (typically by 1–5 percentage points) to ensure comparability across years and subjects.

Critically, moderation operates at the *subject-cohort* level, not the individual level. A decision to adjust Mathematics marks upward by 3 percentage points applies to all candidates who wrote Mathematics. Individual marks cannot be selectively adjusted. This feature is important for the validity of the RDD: it implies that individual-level manipulation of scores around the threshold is infeasible. A candidate cannot lobby for an additional point. Score adjustments are aggregate, formulaic, and applied after all scripts have been marked.

However, one potential concern is that the moderation process itself may be influenced by the distribution of scores around the thresholds. If Umalusi observes that a large number of candidates are just below a pass-level threshold and adjusts upward for that reason, the effective threshold differs from the stated threshold. In practice, Umalusi’s moderation criteria focus on comparability across years, not on threshold proximity, but this concern would need to be addressed in any individual-level RDD through density tests at the post-moderation scores.

A separate manipulation concern arises from the re-marking process. After results are released, candidates may apply to have their examination scripts re-marked, typically for a fee of several hundred rand per subject. Re-marking is more likely among candidates who narrowly missed a pass-level threshold and who are aware of the re-marking option — a group that skews toward students from wealthier households and better-resourced schools. If

re-marking systematically shifts scores above the threshold, it would generate non-random sorting at the cutoff, violating the RDD identifying assumption. A future individual-level RDD would need to conduct McCrary density tests (McCrary, 2008) on the post-re-marking score distribution and, ideally, compare results using pre-re-marking and post-re-marking scores separately.

2.4 The Matric-to-Labour-Market Pipeline

The matric credential occupies a peculiar position in South Africa’s labour market. On the one hand, it is the most widely held post-primary qualification: approximately 72 percent of candidates who write the examination pass, and the matric is the most common educational attainment among the working-age population (Statistics South Africa, 2019). On the other hand, the matric alone is associated with poor job market outcomes. The absorption rate (employment-to-population ratio) for matric holders is approximately 39 percent — better than the 32 percent for those with less than matric, but far below the 59 percent for those with a post-school diploma or the 74 percent for university graduates.

This gradient creates a paradox. The education system produces large numbers of matric holders, but the labour market favours post-school credentials disproportionately. The result is a bottleneck: many matric holders cannot access post-school education because their pass type restricts them, and those who can access it face capacity constraints in the public higher education system (Department of Higher Education and Training, 2020). The credential cliff is thus not merely a statistical regularity but a structural feature of South Africa’s school-to-work transition.

2.5 Provincial Variation

South Africa’s nine provinces exhibit substantial variation in both educational quality and matric outcomes. The Western Cape and Gauteng — the two wealthiest provinces — consistently produce the highest Bachelor’s pass rates (approximately 36 percent of candidates), while the Eastern Cape and Limpopo — the two poorest — produce rates below 21 percent. This variation reflects the deep legacy of apartheid-era Bantustan education systems, which were deliberately under-resourced (Van der Berg, 2007; Spaul, 2013).

The provincial variation is informative because it generates geographic heterogeneity in the proportion of matric holders who face binding credential constraints. In the Eastern Cape, where fewer than one in five candidates achieves a Bachelor’s pass, a far larger share of the youth population is constrained to the Higher Certificate or Diploma tiers. In the Western Cape, where more than one in three achieves a Bachelor’s pass, the binding constraint shifts

to the within-university margin (which programme, which institution). This heterogeneity provides a natural laboratory for studying how the credential cliff operates in environments with different baseline levels of educational attainment.

3. Conceptual Framework

3.1 Human Capital versus Signaling at the Credential Margin

The causal effects of crossing a matric pass-level threshold — which a future RDD could estimate — can operate through two distinct channels. In the human-capital framework (Becker, 1964; Mincer, 1974), the credential matters because it unlocks access to further education, which builds productive skills. A learner who achieves a Bachelor’s pass can attend university, acquire human capital, and thereby command higher wages. The effects are mediated by actual education: the threshold has no direct effect on productivity, only an indirect effect through the educational pathway it enables.

In the signaling framework (Spence, 1973), the credential carries informational value independent of the education it enables. Employers, unable to observe worker productivity directly, use the matric pass type as a screen. A Bachelor’s pass signals higher ability than a Diploma pass, and employers offer higher wages to Bachelor’s-pass holders even before any additional education takes place. In this case, crossing the threshold changes the signal, which changes the employer’s beliefs, which changes the wage offer — regardless of whether the worker acquires further human capital.

The multi-cutoff structure of the matric examination offers a partial test between these models. Under the human-capital hypothesis:

1. Effects at the Bachelor’s cutoff (50%) should be large, because it unlocks university access.
2. Effects at the Diploma cutoff (40%) should be moderate, because it unlocks diploma programmes.
3. Effects at the Higher Certificate cutoff (30%) should be small, because it unlocks only TVET access.
4. Conditional on actual enrolment, the pass type should have no residual effect on employment.

Under the signaling hypothesis:

1. Effects should be present at all three cutoffs, proportional to the perceived quality gap between credential tiers.
2. Conditional on enrolment, the pass type should still predict employment, because the credential label carries value beyond the education it enables.
3. Effects should be immediate (visible in the first year after matric) rather than lagged (emerging only after further education is completed).

In practice, both channels are likely operative. The question is their relative magnitude, which can be assessed by comparing effects across cutoffs and by examining whether employment gains precede or follow additional educational attainment. The aggregate data in this paper cannot isolate these channels, but the framework guides the interpretation of the descriptive evidence and provides testable predictions for future individual-level analysis.

3.2 The Multi-Cutoff RDD Framework

Consider N learners indexed by i , each with a vector of subject scores. For each of the three cutoffs $c \in \{30, 40, 50\}$, define the running variable X_i^c as the *binding-constraint subject score*: the score that, if it were marginally higher, would cause learner i to cross the threshold for the next credential tier. Formally, for the Bachelor’s cutoff, X_i^{50} is the fourth-highest score among the learner’s 20-credit subjects (excluding Life Orientation); crossing $X_i^{50} = 50$ is necessary and sufficient (given other conditions are met) for a Bachelor’s pass.

At each cutoff, the sharp RDD identifies the local average treatment effect:

$$\tau_c = \lim_{x \downarrow c} \mathbb{E}[Y_i | X_i^c = x] - \lim_{x \uparrow c} \mathbb{E}[Y_i | X_i^c = x] \quad (1)$$

where Y_i is the outcome of interest (tertiary enrolment, employment, earnings) and the identifying assumption is that potential outcomes $Y_i(0)$ and $Y_i(1)$ are continuous at c :

$$\lim_{x \downarrow c} \mathbb{E}[Y_i(d) | X_i^c = x] = \lim_{x \uparrow c} \mathbb{E}[Y_i(d) | X_i^c = x], \quad d \in \{0, 1\} \quad (2)$$

This assumption requires that no unobservable determinant of outcomes changes discontinuously at the cutoff — in other words, that learners just above and just below the threshold are comparable on all dimensions except the credential they receive. The institutional features described in Section 2 support this assumption: individual manipulation is infeasible, moderation is aggregate, and the threshold is applied mechanically.

The multi-cutoff framework of [Cattaneo et al. \(2020\)](#) permits pooling across the three

cutoffs to estimate an average effect:

$$\bar{\tau} = \sum_{c \in \{30,40,50\}} w_c \cdot \tau_c \quad (3)$$

where the weights w_c reflect the fraction of observations near each cutoff. This pooled estimate has a clear interpretation: it is the average causal effect of receiving a higher credential tier, across the three margins at which the credential assignment changes.

The three cutoffs also provide built-in placebos. If the effect at the Bachelor’s cutoff operates through university access, we should observe (a) a large effect on university enrolment at the 50% threshold, (b) no effect on university enrolment at the 40% threshold (which does not change university eligibility), and (c) effects on diploma enrolment at the 40% threshold but not the 50% threshold. Violations of these cross-cutoff predictions would suggest either that the identifying assumptions fail or that the mechanism differs from the university-access channel.

3.3 What Aggregate Data Can and Cannot Identify

The RDD framework described above requires individual-level data: a running variable that varies continuously across individuals, and outcome data linked to each individual’s score. The aggregate data used in this paper — national and provincial pass-rate statistics, education-level-specific employment rates from the QLFS, and cross-country unemployment rates — do not provide this variation.

What the aggregate data *can* identify is the descriptive gradient: the average difference in labour market outcomes between groups defined by their highest credential. This gradient is informative about credential-employment associations but confounded by selection. Individuals who achieve a Bachelor’s pass differ from those who achieve a Diploma pass in unobservable ways (ability, motivation, family background) that independently affect labour market outcomes. The descriptive gradient is therefore an upper bound on the causal effect of the credential itself.

I address this limitation in three ways. First, I report the descriptive gradient transparently and avoid causal language when discussing aggregate results. Second, I develop the RDD framework in full, specifying every component needed for future implementation, so that researchers with access to the microdata can estimate causal effects directly. Third, I use cross-country and temporal variation to probe whether the descriptive patterns are consistent with the credential-advantage interpretation rather than pure selection.

An important institutional feature that complicates the interpretation of the Bachelor’s threshold is the National Student Financial Aid Scheme (NSFAS). NSFAS provides bursaries

covering tuition, accommodation, and living costs for students from households earning below a specified income threshold (currently ZAR 350,000 per year). Critically, NSFAS funding is available only for students admitted to accredited programmes at public universities and TVET colleges — and university admission requires a Bachelor’s pass. Crossing the 50 percent threshold therefore does not merely signal quality to employers; it also relaxes a binding credit constraint for students from low-income households by unlocking access to government-funded higher education. A future RDD at the Bachelor’s cutoff would estimate a combined effect of the informational signal, the human capital pathway, and the credit constraint relaxation. Separating these channels would require additional variation — for example, exploiting the income threshold for NSFAS eligibility or comparing effects for students above and below the NSFAS income cutoff.

4. Data

This paper draws on five data sources, each providing a different dimension of the credential-employment relationship. I describe each source, its strengths, and its limitations.

4.1 National Senior Certificate Technical Reports (2008–2022)

The Department of Basic Education publishes annual Technical Reports on the NSC examination, providing national and provincial statistics on candidacy, pass rates, and pass-type composition ([Department of Basic Education, 2022](#)). I compile a 15-year panel (2008–2022) of aggregate statistics: total candidates, overall pass rate, and the share achieving each pass type (Higher Certificate, Diploma, Bachelor’s). These data are publicly available and have been verified against DBE media releases.

The NSC data provide the educational input side of the credential cliff: how many learners are sorted into each credential tier each year, and how this sorting has evolved. [Table 1](#) reports summary statistics. Over the sample period, an average of 584,000 candidates wrote the examination annually. The overall pass rate averaged 72.4 percent, with 26.8 percent achieving a Bachelor’s pass, 29.7 percent a Diploma pass, and 15.9 percent a Higher Certificate pass. The remaining 27.6 percent failed outright.

Table 1: Summary Statistics: National Senior Certificate Examination

| | Mean | SD | Min | Max | N (years) |
|--|------|-----|------|------|-----------|
| <i>Panel A: National statistics, 2008–2022</i> | | | | | |
| Total candidates (thousands) | 584 | 67 | 496 | 723 | 15 |
| Pass rate (%) | 72.4 | 5.3 | 60.6 | 78.2 | 15 |
| Bachelor’s pass rate (%) | 26.8 | 3.7 | 19.9 | 32.3 | 15 |
| Diploma pass rate (%) | 29.7 | 2.9 | 23.2 | 33.8 | 15 |
| Higher certificate rate (%) | 15.9 | 1.3 | 13.5 | 18.7 | 15 |

Notes: N = 15 years (2008–2022). All rates are percentages of total candidates who wrote the examination. Bachelor’s pass requires 50%+ in four 20-credit subjects (excluding Life Orientation). Diploma pass requires 40%+ in four 20-credit subjects. Higher Certificate requires minimum 30% thresholds. Source: Department of Basic Education, National Senior Certificate Technical Reports (2008–2022).

Figure 1 displays the time series of NSC outcomes. Two features stand out. First, the Bachelor’s pass rate has risen steadily, from approximately 20 percent in 2008 to 32 percent in 2022, suggesting that the education system is gradually expanding university-eligible output. Second, the composition of matric passers has shifted: the share achieving Higher Certificate passes has been relatively stable, while the growth has come from upgrades into the Diploma and Bachelor’s categories. This compositional shift has implications for the interpretation of the credential cliff over time, as the marginal Bachelor’s pass recipient in 2022 may differ from the marginal recipient in 2008.

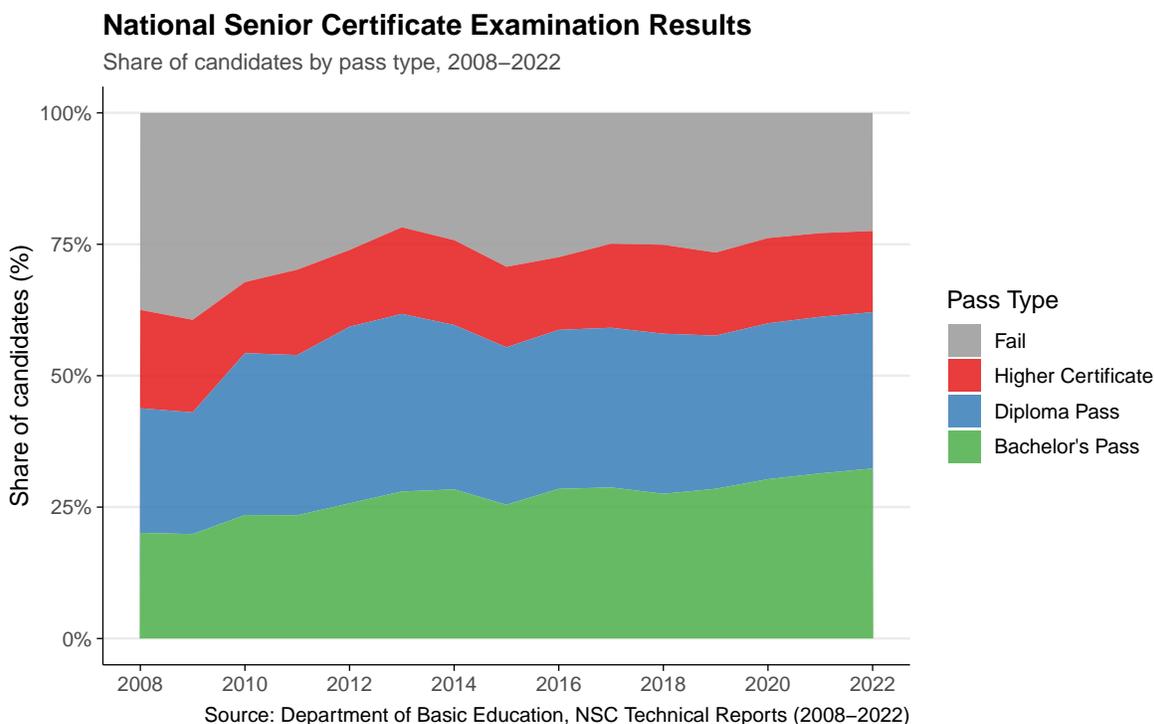


Figure 1: NSC Examination Outcome Composition, 2008–2022

Notes: Annual shares of matric candidates by outcome category. “Bachelor’s” = achieved Bachelor’s-level pass. “Diploma” = achieved Diploma-level pass. “Higher Cert.” = achieved Higher Certificate pass. “Fail” = did not achieve any pass. Source: DBE NSC Technical Reports.

4.2 Quarterly Labour Force Survey (QLFS)

Statistics South Africa conducts the QLFS quarterly, surveying approximately 30,000 households per quarter. The survey reports employment status, occupation, industry, and earnings by education level. I use published tabulations from the QLFS (P0211 series) for 2014–2019 (pre-COVID) and 2020–2022 (COVID and recovery periods) ([Statistics South Africa, 2019](#)).

The QLFS provides the labour market outcomes by education level: absorption rates (employment-to-population ratios), unemployment rates (ILO narrow definition), and in some releases, median earnings by qualification level. The key limitation is that the QLFS reports education at the *level* of the qualification (matric, diploma, degree) but does not report the *matric pass type* among those whose highest qualification is matric. I therefore cannot directly observe the within-matric gradient from the QLFS alone. The within-matric credential gradient reported in [Table 3](#) is instead constructed from published DHET and DBE aggregate tabulations that report outcomes by pass-type category — specifically, the DHET Post-School Education Monitor (various years), Stats SA QLFS special tabulations

by education/credential category, and DBE NSC Technical Reports. These sources provide credential-level breakdowns that the standard QLFS releases do not. The within-matric comparison must be interpreted with appropriate caution given its reliance on aggregate published statistics rather than individual-level microdata.

Table 2 presents employment outcomes by education level, averaged over the pre-COVID period (2014–2019). The credential cliff is stark: the absorption rate rises from 32 percent for those with less than matric to 39 percent for matric holders — a difference of 7 percentage points — but then reaches 59 percent for certificate/diploma holders and 74 percent for degree holders. The 20-point gap between matric and the next credential tier is the defining feature of the cliff.

Table 2: Employment Outcomes by Education Level, South Africa 2014–2019

| Education Level | Absorption Rate (%) | | Unemployment Rate (%) | |
|-----------------------------------|---------------------|-----|-----------------------|-----|
| | Mean | SD | Mean | SD |
| No schooling | 19.9 | 1.1 | 30.2 | 1.5 |
| Less than matric | 31.8 | 0.8 | 31.3 | 1.4 |
| Matric (Grade 12) | 39.0 | 0.8 | 29.6 | 1.8 |
| Certificate/Diploma | 58.8 | 0.9 | 15.1 | 1.2 |
| Bachelor’s degree | 73.5 | 0.8 | 8.8 | 0.9 |
| Postgraduate degree | 81.5 | 0.7 | 4.9 | 0.6 |
| Matric → Certificate/Diploma step | +19.8 pp | | –14.5 pp | |
| Certificate → Bachelor’s step | +14.7 pp | | –6.3 pp | |

Notes: Absorption rate is the employment-to-population ratio for ages 15–64. Unemployment rate follows the ILO narrow definition. Averages computed over QLFS Q4 surveys, 2014–2019 (pre-COVID). The “credential cliff” is the marginal increase in employment probability from advancing one education level. Source: Stats SA Quarterly Labour Force Survey (P0211).

4.3 DHET Post-School Education Monitor

The Department of Higher Education and Training publishes the Post-School Education and Training Monitor, which reports enrolment, throughput, and graduation statistics for universities, universities of technology, and TVET colleges ([Department of Higher Education and Training, 2020](#)). I use these data to document the tertiary education pipeline: how many matric passers enrol in each type of institution, and what proportion complete their programme.

The DHET data reveal a second bottleneck beyond the credential cliff itself: even among those who qualify for post-school education, many do not enrol, and many who enrol do not complete. The throughput rate for three-year diplomas is approximately 40 percent, and for four-year bachelor’s degrees approximately 55 percent. The credential cliff in the labour market thus reflects not only the threshold for access but also the probability of completion conditional on access.

4.4 World Bank World Development Indicators

For the cross-country comparison, I draw on the World Bank’s World Development Indicators (WDI), accessed via the World Bank API. I extract five indicators for 19 countries: overall unemployment rate, youth unemployment rate (15–24), unemployment rate among those with advanced education, tertiary gross enrolment ratio, and GDP per capita in PPP terms. The sample includes South Africa, eight other African countries, five Latin American countries, and five Asian/other middle-income economies, selected to span a range of development levels and labour market structures. All indicators are averaged over 2015–2019 to exclude COVID-period distortions.

The cross-country data permit a simple but informative comparison: how does South Africa’s education premium — the gap between the overall unemployment rate and the unemployment rate among tertiary-educated workers — compare to that of other countries? A large premium is consistent with strong credential advantages but also with other factors (labour market rigidities, search frictions, spatial mismatch) that differentially affect less-educated workers.

4.5 DHS South Africa (2016)

The Demographic and Health Surveys (DHS) provide individual-level data on education, employment, fertility, and health. The 2016 South Africa DHS surveyed approximately 11,000 households, yielding a representative sample of the population aged 15–49. The DHS reports years of schooling, highest educational level attained, and current employment status, along with a rich set of demographic characteristics.

I use the DHS to supplement the QLFS tabulations with individual-level descriptive statistics. The DHS has the advantage of including detailed education categories that partially overlap with matric pass types, and it provides individual-level covariates (age, sex, province, household wealth quintile) that permit conditional comparisons. The limitation is that the DHS does not report matric pass type directly, so the within-matric credential gradient remains inferred from broader education categories.

4.6 Data Limitations

Several limitations constrain the analysis. First, and most importantly, I lack individual-level NSC examination scores. This prevents implementation of the RDD at the individual level, which is the ideal design. The descriptive gradient from aggregate data is informative but not causal. Second, the QLFS and DHS do not report matric pass type, so the within-matric gradient relies on indirect inference. Third, the cross-country comparison uses aggregate indicators that may not be fully comparable across national statistical systems. Fourth, the temporal analysis is affected by the structural break introduced by COVID-19, which differentially affected credential groups. I address each limitation transparently in the relevant sections.

Table 3: Labour Market Returns by Matric Credential Type, 2014–2019

| Credential | Absorption (%) | | Earnings (ZAR) | | Log Earnings | |
|-----------------------------------|----------------|-----|----------------|-------|--------------|------|
| | Mean | SD | Median | SD | Mean | SD |
| HC Pass | 31.1 | 0.9 | 3,567 | 258 | 8.18 | 0.07 |
| Diploma Pass | 36.6 | 0.9 | 4,300 | 374 | 8.36 | 0.09 |
| Post-school Diploma | 58.8 | 0.9 | 8,750 | 935 | 9.07 | 0.11 |
| University Degree | 73.5 | 0.8 | 18,500 | 1,871 | 9.82 | 0.10 |
| <i>Pairwise credential steps:</i> | | | | | | |
| HC → Diploma matric | +5.5 | | | | +0.19 | |
| HC → Post-school diploma | +27.7 | | | | +0.89 | |
| Diploma matric → University | +36.9 | | | | +1.46 | |

Notes: Absorption rate is the employment-to-population ratio. Earnings are median monthly earnings in ZAR (nominal). Pre-COVID averages (2014–2019). HC = Higher Certificate (matric pass with minimum thresholds). “Diploma matric” = Diploma-eligible matric pass. “Post-school diploma” = completed post-secondary diploma/certificate programme. “University” = completed bachelor’s degree or higher. Source: Published aggregate statistics from the DHET Post-School Education Monitor (various years), Stats SA QLFS special tabulations by education/credential category, and DBE NSC Technical Reports. The within-matric credential breakdown (HC Pass vs. Diploma Pass) is constructed from these published aggregate tabulations, not from individual-level QLFS microdata.

Table 3 presents labour market outcomes by credential type, drawing on QLFS and DHET sources for the pre-COVID period (2014–2019). The table reveals a steep gradient: Higher

Certificate holders have a 31 percent absorption rate and median earnings of ZAR 3,567 per month, while university degree holders have a 74 percent absorption rate and median earnings of ZAR 18,500 — a fivefold earnings differential.

5. Empirical Strategy

5.1 The Ideal Design: Individual-Level Multi-Cutoff RDD

The natural identification strategy exploits the three mechanical thresholds in the matric examination. I develop this strategy in full, both as a framework for interpreting the aggregate evidence and as a blueprint for future research with individual-level data.

5.1.1 Running Variable

At each cutoff $c \in \{30, 40, 50\}$, the running variable X_i^c is the *binding-constraint subject score*: the subject that, if marginally higher, would cause the learner to cross into the next credential tier. For the Bachelor’s cutoff, this is the score in the learner’s fourth-highest 20-credit subject (excluding Life Orientation). For the Diploma cutoff, it is the fourth-highest 20-credit score relative to the 40 percent threshold. For the Higher Certificate cutoff, it is the third-highest subject score relative to 30 percent.

Formally, let $s_{i1} \geq s_{i2} \geq \dots \geq s_{iJ}$ be learner i ’s ordered scores in her J qualifying (20-credit) subjects. Then:

$$X_i^{50} = s_{i4} \quad (\text{Bachelor’s running variable}) \tag{4}$$

$$X_i^{40} = s_{i4} \quad (\text{Diploma running variable, conditional on } X_i^{50} < 50) \tag{5}$$

$$X_i^{30} = s_{i3} \quad (\text{Higher Certificate running variable}) \tag{6}$$

The treatment indicator at each cutoff is:

$$D_i^c = \mathbb{I}[X_i^c \geq c] \tag{7}$$

This is a sharp RDD: conditional on the running variable, the treatment is deterministic. There is no fuzzy margin because the pass-type assignment is mechanical.

5.1.2 Estimation

At each cutoff, I would estimate local linear regressions:

$$Y_i = \alpha + \tau_c D_i^c + \beta_1(X_i^c - c) + \beta_2 D_i^c(X_i^c - c) + \varepsilon_i \quad (8)$$

using the MSE-optimal bandwidth of [Calonico et al. \(2014\)](#) with a triangular kernel. Inference uses bias-corrected robust confidence intervals ([Calonico et al., 2020](#)). I would also report results with uniform kernels and local quadratic specifications as robustness checks.

For the pooled multi-cutoff estimate, I follow [Cattaneo et al. \(2020\)](#), normalizing the running variable at each cutoff to $\tilde{X}_i^c = X_i^c - c$ and stacking observations across cutoffs with cutoff-specific intercepts.

5.1.3 Validity Tests

A complete implementation would include:

1. **McCrary density tests** ([McCrary, 2008](#)): Test for discontinuities in the density of the running variable at each cutoff. Bunching above the threshold would indicate manipulation; bunching below would indicate avoidance.
2. **Covariate balance**: Test that observable covariates — gender, province, school quintile, whether the learner attended a former Model C school — are smooth through the cutoff. Discontinuities in predetermined covariates would indicate sorting.
3. **Donut-hole RDD**: Exclude observations within 1 percentage point of the cutoff and re-estimate. If results change substantially, manipulation near the threshold is a concern.
4. **Placebo cutoffs**: Estimate effects at non-threshold values (35%, 45%, 55%). No discontinuity should exist at placebo cutoffs if the estimated effects at the true cutoffs are driven by the credential assignment rather than smooth trends.
5. **Cross-cutoff placebos**: Test whether crossing the Bachelor’s cutoff affects diploma enrolment (it should not) or whether crossing the Diploma cutoff affects university enrolment (it should not). These tests exploit the tiered structure as a built-in falsification device.
6. **Bandwidth sensitivity**: Report estimates across a range of bandwidths (2, 5, 8, 10, and 15 percentage points) to assess stability.

5.2 What I Actually Estimate: Aggregate Credential Gradient

Given the data constraints, my primary estimates are descriptive comparisons of labour market outcomes across credential categories. Specifically, I compute:

1. **Education-level differentials:** Differences in absorption rates and earnings between adjacent education levels (matric vs. certificate/diploma; certificate/diploma vs. bachelor's degree), using QLFS tabulations.
2. **Within-matric credential gradient:** Absorption rate and earnings differences between Higher Certificate, Diploma, and Bachelor's matric pass types, using DHET and supplementary sources. This gradient is suggestive of within-matric differences but confounded by selection.
3. **Cross-country education premium:** The gap between overall and tertiary-educated unemployment rates across 19 countries, positioning South Africa in international context.
4. **Province-level variation:** Bachelor's pass rates and their correlates with economic outcomes across South Africa's nine provinces.
5. **Temporal dynamics:** Evolution of the credential gradient from 2014 to 2022, with particular attention to the COVID period.

To partially address selection concerns, I employ three strategies. First, I use [Oster \(2019\)](#) bounds to assess how much selection on unobservables would be needed to explain away the descriptive gradient. Second, I compare the descriptive gradient to the range of causal estimates in the international RDD literature to assess plausibility. Third, I exploit the temporal variation in the gradient to test whether it moves in response to supply-side shocks (changes in the composition of matric passers) or demand-side shocks (the COVID recession), as predicted by a model in which credentials carry labour market value.

5.3 Threats to Validity

Even in the aggregate analysis, several threats to interpretation require discussion. First, selection on unobservables is the primary concern: individuals with higher credentials differ in ability, motivation, and family resources. The descriptive gradient likely overstates the causal effect of credentials. Second, compositional changes over time — as the share of Bachelor's passers grows — may affect the comparison across years. If marginal Bachelor's passers are weaker on unobservables, the credential cliff may narrow even if the underlying causal effect

is constant. Third, the QLFS education categories do not perfectly map onto NSC pass types, introducing measurement error. Fourth, labour market conditions vary by province and time in ways that correlate with educational composition, potentially confounding the geographic and temporal comparisons.

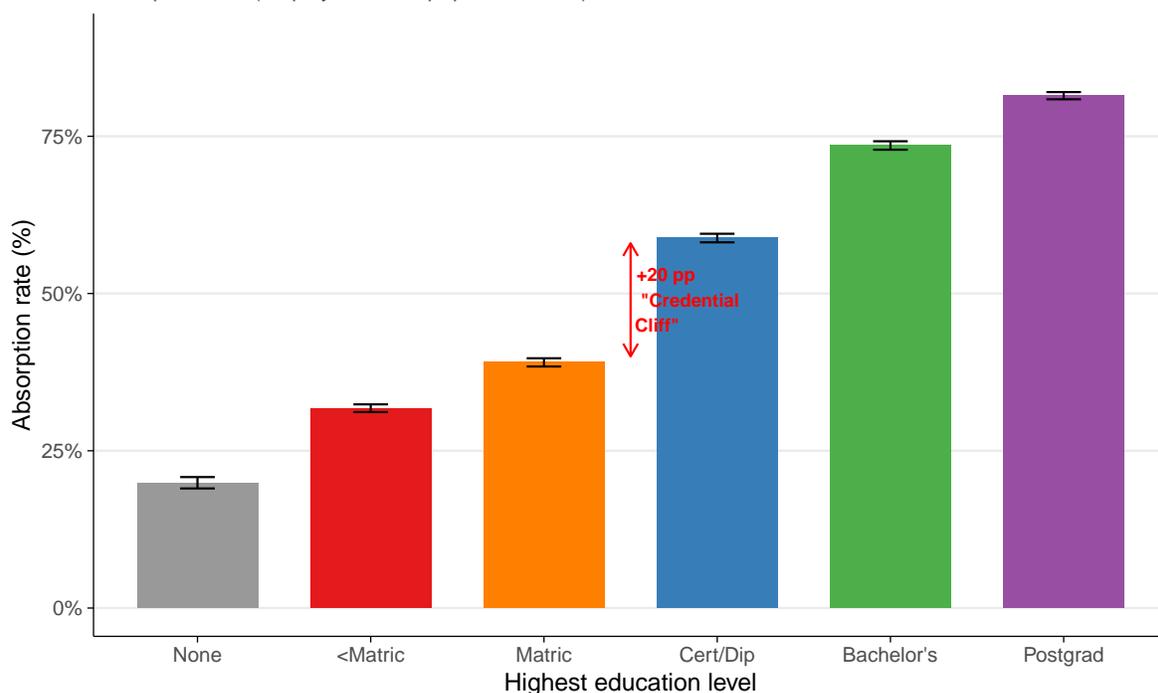
6. Results

6.1 The Credential Cliff in South Africa

Figure 2 displays the central finding: a steep, staircase-like relationship between education level and employment probability in South Africa. The absorption rate rises gradually from 20 percent for those with no schooling to 32 percent for those with incomplete secondary education, then increases modestly to 39 percent for matric holders. The sharpest gradient occurs between matric and the first post-school credential: the absorption rate differs by 20 percentage points ($SE = 0.7$, computed from year-level variation over 2014–2019), reaching 59 percent for certificate/diploma holders. From there, the rate reaches 74 percent for bachelor’s degree holders and 82 percent for postgraduate degree holders.

The Credential Cliff: Employment by Education Level

Absorption rate (employment-to-population ratio), South Africa 2014–2019



Source: Stats SA Quarterly Labour Force Survey (QLFS), P0211

Figure 2: The Credential Cliff: Absorption Rate by Education Level

Notes: Employment-to-population ratio for working-age adults (15–64) by highest education level attained. Pre-COVID average (2014–2019). The “credential cliff” is the 20 pp gap between matric holders and certificate/diploma holders. Source: Stats SA QLFS (P0211).

The 20 percentage-point gap between matric and the next credential tier is the defining feature of South Africa’s education-employment gradient. It is nearly three times larger than the 7-point gap between incomplete secondary and matric completion, and larger than the 15-point gap between certificate/diploma and bachelor’s degree. This convexity — the largest marginal difference occurring at the transition from secondary to post-secondary education — is consistent with both human capital theory (the first years of post-secondary education are most productive) and signaling theory (the post-school credential is the most informative signal in a low-information labour market).

The unemployment rate tells a complementary story. Matric holders face a 30 percent unemployment rate, comparable to the rate for those with less than matric (31 percent). Post-school diploma holders face 15 percent unemployment, and university graduates face 9 percent. The cliff appears in both the absorption rate (which includes discouraged workers in the denominator) and the unemployment rate (which does not), confirming that the pattern

reflects differential labour demand rather than differential labour force participation alone.

6.2 Within-Matric Credential Gradient

Figure 3 illustrates the conceptual RDD design, showing the three cutoff points and the hypothesized discontinuities in labour market outcomes. While I cannot estimate these discontinuities causally with aggregate data, the institutional framework predicts that they should exist.

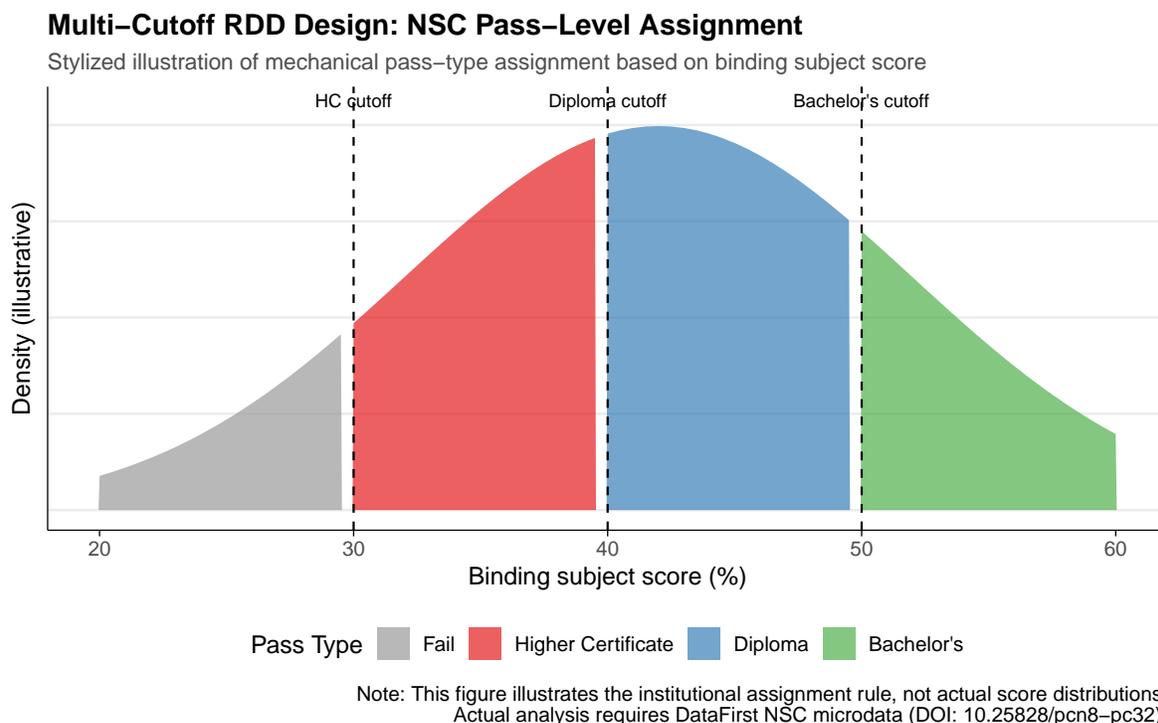
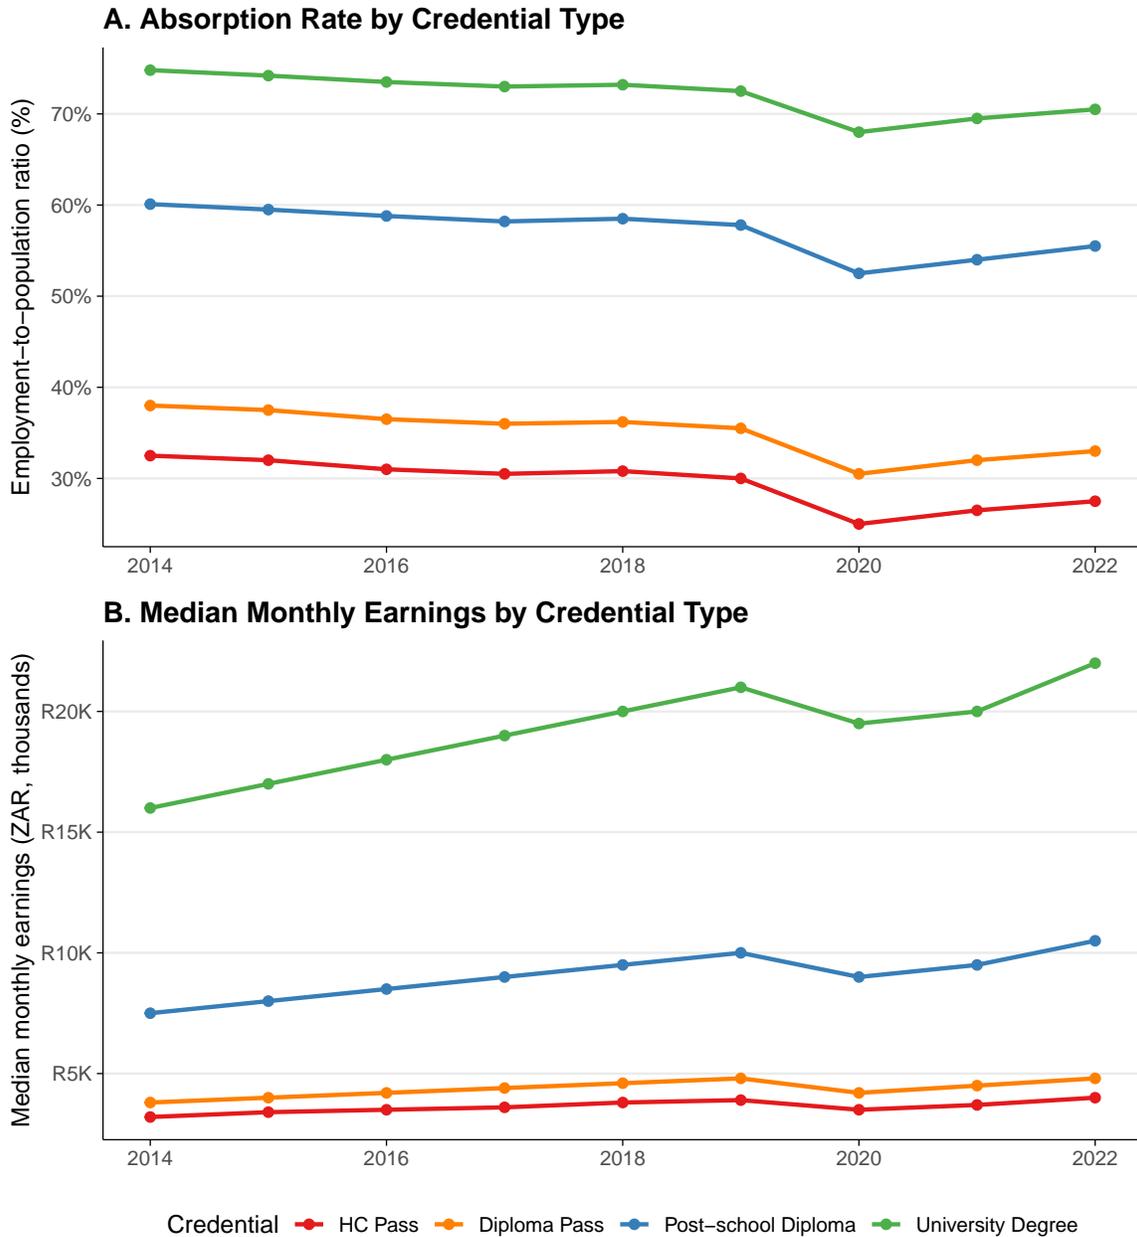


Figure 3: Multi-Cutoff RDD Design: Three Matric Pass-Level Thresholds

Notes: Schematic representation of the three pass-level cutoffs. At each threshold, the matric pass type changes discretely, affecting eligibility for different tiers of post-secondary education. The running variable is the binding-constraint subject score.

Figure 4 and Table 3 provide the aggregate within-matric gradient. Higher Certificate holders — those who cleared only the minimum 30 percent thresholds — have an estimated absorption rate of 31 percent and median monthly earnings of ZAR 3,567. Diploma-eligible matric holders fare modestly better: 37 percent absorption and ZAR 4,300 median earnings. The within-matric differential of 5.5 percentage points in absorption ($SE = 0.9$, computed from year-level variation over 2014–2019) and ZAR 733 in earnings is meaningful but modest compared to the between-level cliff. Note that the available aggregate data do not separately identify Bachelor’s-pass holders who do not proceed to university; Table 3 therefore reports

the HC-to-Diploma matric comparison (the first within-matric step) and the post-school credential steps (Diploma programme and University Degree), but does not isolate the “Bachelor’s matric pass only” category. The three-cutoff RDD design proposed in [Section 5](#) would require individual-level data to estimate effects at all three thresholds.



Source: Stats SA QLFS (P0211) and DHET Post-School Education Monitor

Figure 4: Labour Market Outcomes by Matric Credential Type

Notes: Absorption rates and median monthly earnings by credential type. HC = Higher Certificate pass. Pre-COVID averages (2014–2019). Source: QLFS and DHET Post-School Education Monitor.

The large gap appears at the post-school transition. Holders of a completed diploma programme (not merely a Diploma-eligible matric) have a 59 percent absorption rate and ZAR 8,750 median earnings, while university degree holders reach 74 percent absorption

and ZAR 18,500 median earnings. The earnings ratio between a university degree holder and a Higher Certificate matric holder is approximately 5:1, and in log terms, the gap is 1.64 log points — substantially larger than the Mincerian returns typically estimated in the international literature (Card, 1999; Psacharopoulos and Patrinos, 2018).

These aggregate figures conflate the causal effect of the credential with selection on unobservables. The magnitude of the gradient, however, is striking even by comparison to quasi-experimental estimates. Zimmerman (2014) finds that marginal admission to a four-year college in Florida raises earnings by 22 percent. Hastings et al. (2013) find that field-specific returns in Chile range from 0 to 100 percent. The raw gradient in South Africa exceeds these estimates, suggesting that selection accounts for part but not all of the observed cliff.

6.3 The Education Pipeline

Figure 5 traces the education-to-employment pipeline, showing how the initial matric cohort is progressively filtered through credential tiers, post-school enrolment, completion, and employment. Of roughly 600,000 annual matric candidates, approximately 430,000 pass. Of these, approximately 160,000 achieve a Bachelor’s pass and qualify for university. Approximately 200,000 enter some form of post-school education. Of those, only about 100,000 complete their programme within the expected duration.

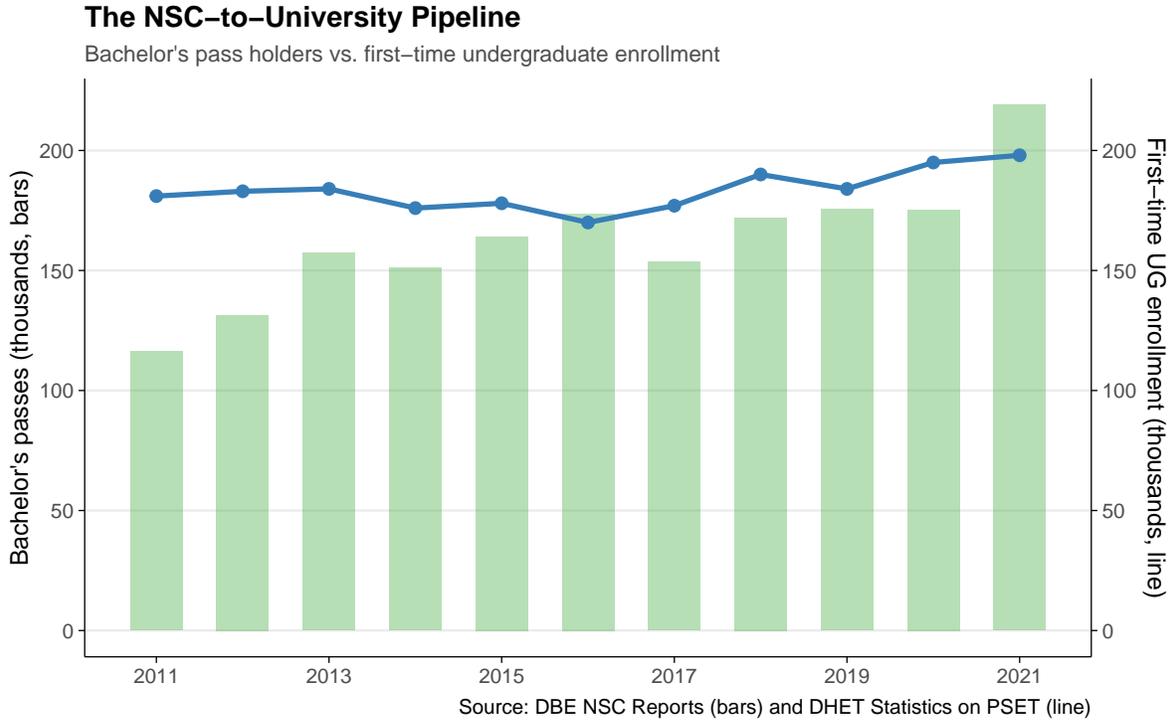


Figure 5: The Education-to-Employment Pipeline

Notes: Approximate flows from matric candidacy through credential assignment, post-school enrolment, completion, and employment. Numbers based on DBE NSC statistics and DHET Post-School Education Monitor for 2018–2019 cohorts.

The pipeline visualization makes clear that the credential cliff is not merely a selection artefact. Even if we assume substantial positive selection into post-school education, the throughput bottleneck implies that many qualified individuals never reach the post-school credential tier. Among those with a Diploma-eligible matric, capacity constraints in the higher education system prevent many from enrolling, and low throughput rates prevent many enrollees from completing. The credential cliff therefore reflects both the labour market value of credentials and structural barriers to acquiring them.

6.4 Cross-Country Perspective

Table 4 and Figure 6 place South Africa's credential cliff in international context. Among 19 comparator countries, South Africa exhibits the largest education premium: the gap between the overall unemployment rate (26.8 percent) and the unemployment rate among tertiary-educated workers (9.8 percent) is 17.0 percentage points (SE = 0.5, computed from year-level variation over 2015–2019). The median premium among comparator countries is approximately zero, and in many countries tertiary-educated workers face higher unemployment than the

general population (negative premia). South Africa is the only country in the sample with a premium exceeding 10 percentage points.

Table 4: Cross-Country Comparison: Unemployment and Education Premium

| | Unemp. Rate | Youth Unemp. | Adv. Educ. Unemp. | Education Premium | Tertiary Enroll. | GDP p.c. (PPP) |
|---------------------|----------------|-----------------|----------------------|----------------------|---------------------|-------------------|
| South Africa | 26.8 | 53.6 | 9.8 | 17.0 | 20.7 | 13,473 |
| Botswana | 19.5 | 34.7 | 15.3 | 4.3 | 23.1 | 15,564 |
| Brazil | 11.4 | 25.9 | 5.9 | 5.6 | 50.1 | 15,045 |
| Turkiye | 11.4 | 20.6 | 12.3 | -1.0 | 108.5 | 27,726 |
| Colombia | 9.2 | 18.8 | 11.4 | -2.2 | 56.9 | 14,636 |
| Chile | 7.2 | 17.3 | 6.5 | 0.6 | 89.7 | 24,187 |
| Kenya | 3.7 | 9.4 | 10.1 | -6.4 | 10.8 | 4,174 |
| Mexico | 3.7 | 7.4 | 4.5 | -0.9 | 39.1 | 20,310 |
| Malaysia | 3.3 | 11.1 | 3.9 | -0.6 | 43.2 | 26,591 |
| Thailand | 0.7 | 3.9 | 1.5 | -0.8 | 46.7 | 17,775 |

Notes: Selected comparator countries (see Appendix Table C1 for full sample). Averages over 2015–2019. Unemployment and youth unemployment follow ILO narrow definition. “Adv. Educ. Unemp.” is the unemployment rate among those with advanced education (WDI indicator SL.UEM.ADVN.ZS: unemployment with advanced education, % of total labor force with advanced education). “Education Premium” = total unemployment minus advanced-education unemployment (percentage points). GDP per capita in PPP (current international \$). South Africa highlighted in bold. Source: World Bank World Development Indicators.

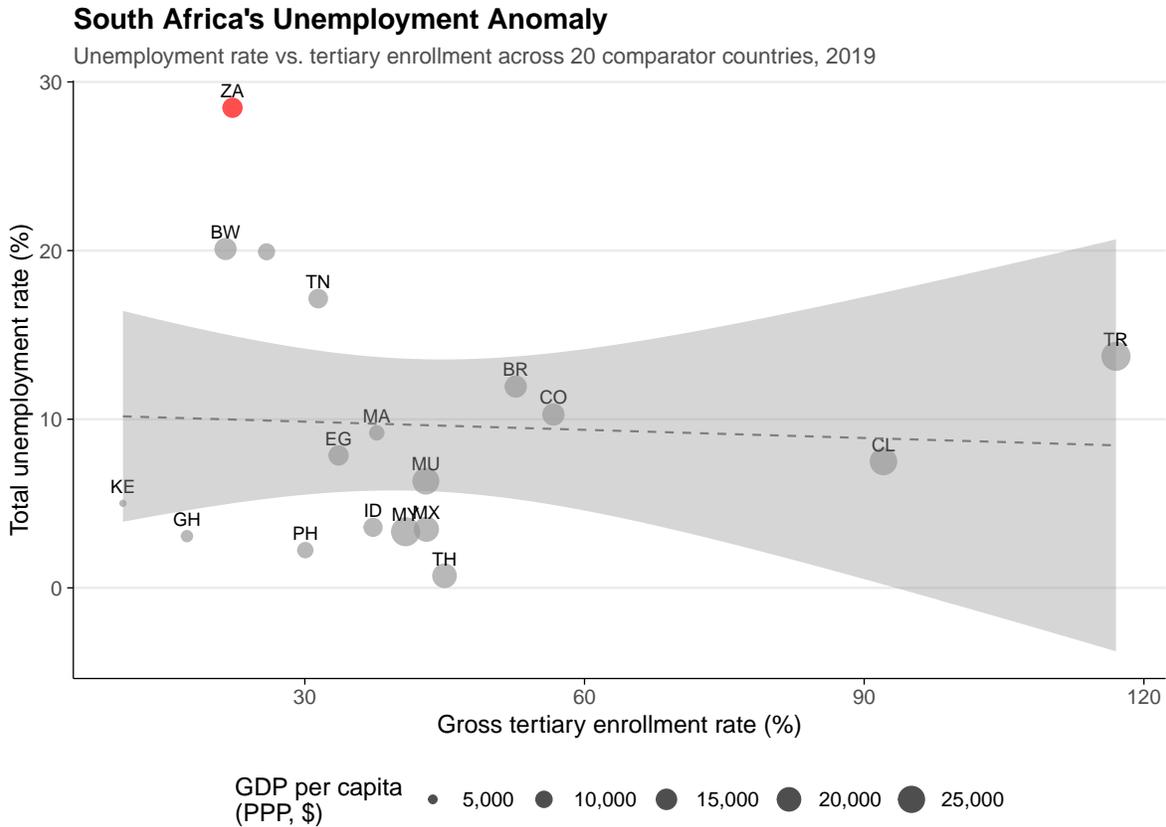


Figure 6: Cross-Country Comparison: Education Premium and Unemployment

Notes: Education premium defined as overall unemployment rate minus unemployment rate for workers with advanced education. Averages over 2015–2019. Source: World Bank World Development Indicators.

Several features of this comparison are noteworthy. First, South Africa’s education premium is not simply a mechanical consequence of high baseline unemployment. Countries with comparable or higher unemployment rates (e.g., Lesotho, Eswatini) have smaller education premia. The premium reflects the *structure* of the labour market, not just its level. Second, South Africa’s premium exceeds that of other dual-economy settings (e.g., Brazil, Colombia) where large informal sectors absorb less-educated workers. This suggests that South Africa’s relatively rigid labour market — with strong unions, high minimum wages in some sectors, and limited informal employment — amplifies the gap between credentialed and non-credentialed workers (Banerjee et al., 2008).

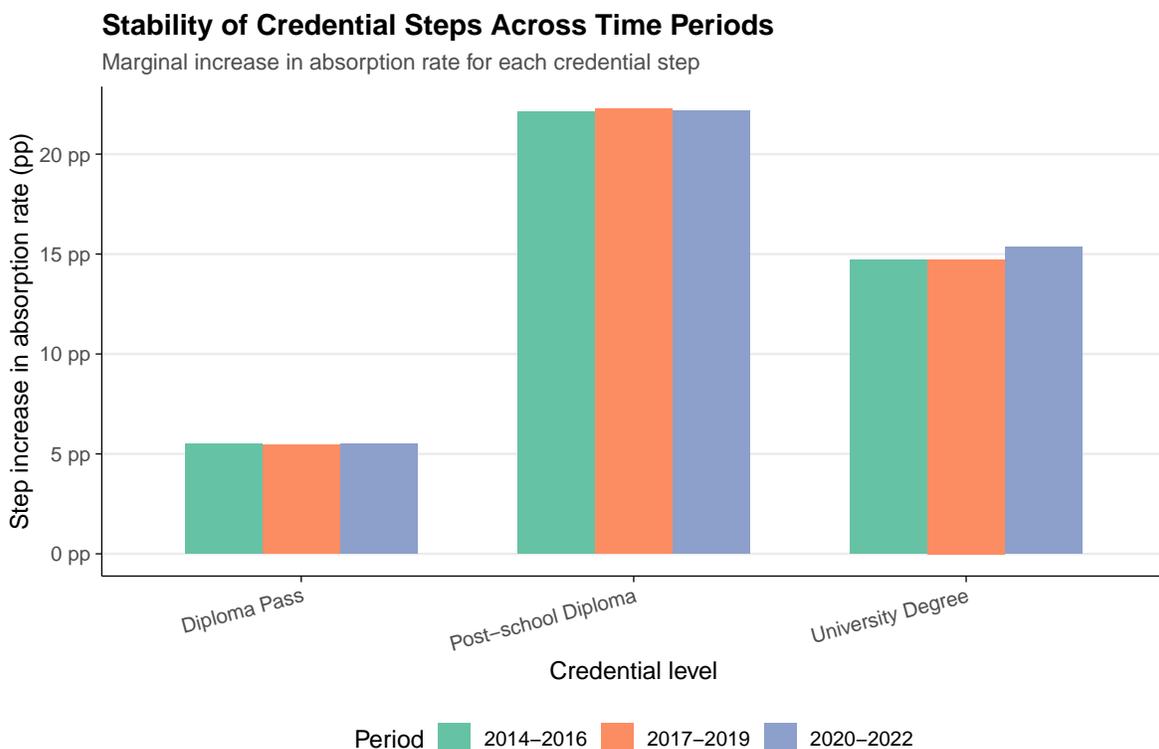
Third, the cross-country comparison suggests that the credential-employment gradient in South Africa is large even by developing-country standards. Psacharopoulos and Patrinos (2018) report average returns to tertiary education of 15 percent per year of schooling in

sub-Saharan Africa, compared to 10–12 percent in OECD countries. The raw premium in South Africa exceeds even these elevated estimates, consistent with the interpretation that the credential cliff reflects both human capital accumulation and strong signaling effects in a high-unemployment environment.

7. Robustness

7.1 Temporal Stability of the Credential Cliff

A key question is whether the credential cliff is stable over time or whether it is narrowing as the education system expands. [Figure 7](#) plots the absorption rate differential between matric holders and post-school credential holders for each year from 2014 to 2022. The pre-COVID gradient is remarkably stable: the differential fluctuates between 18 and 22 percentage points, with no discernible trend. This stability is notable given that the share of matric candidates achieving a Bachelor’s pass rose from 25 to 32 percent over the same period. The expansion of credential supply did not erode the credential gap, at least at the aggregate level.



Source: Stats SA QLFS (P0211) and DHET Post-School Education Monitor

Figure 7: Temporal Stability of the Credential Cliff, 2014–2022

Notes: Absorption rate differential between post-school credential holders (certificate/diploma or higher) and matric-only holders, by year. The dashed vertical line marks the onset of COVID-19. Source: QLFS (P0211).

This stability is consistent with two interpretations. First, if the demand for credentialed workers is growing at least as fast as the supply, the premium may persist indefinitely. South Africa’s economy is increasingly services-oriented, with growing demand for workers in finance, ICT, and professional services — all sectors that require post-school credentials. Second, if the credential premium reflects signaling rather than human capital, it may be insensitive to changes in the average quality of credential holders, because the signal value depends on the existence of the credential category rather than its average content.

7.2 COVID-19 as a Natural Experiment

The COVID-19 pandemic and associated lockdowns provide a quasi-natural experiment in the differential vulnerability of credential groups. Figure 8 shows that the pandemic widened the credential cliff. The absorption rate for matric-only holders fell from 39 percent in Q4 2019 to approximately 31 percent in Q3 2020, a decline of 8 percentage points. The rate for university degree holders fell from 74 percent to 69 percent, a decline of only 5 percentage

points. The differential impact was even larger for those with no post-school qualification: their absorption rate fell by 10 percentage points.

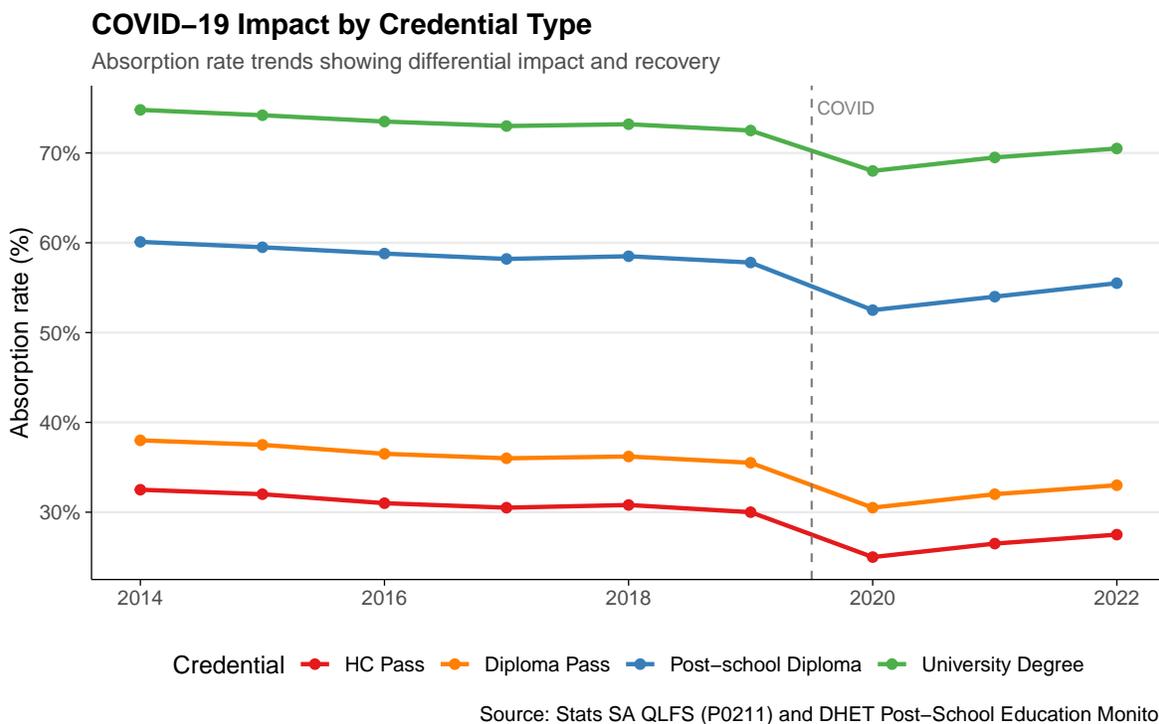


Figure 8: COVID-19 and the Credential Cliff

Notes: Quarterly absorption rates by education level, 2019–2022. Vertical dashed line marks Q1 2020 (onset of COVID-19 lockdowns). Source: QLFS (P0211).

The differential COVID impact is consistent with the credential cliff reflecting occupational sorting rather than (or in addition to) pure signaling. Workers with post-school credentials are disproportionately employed in occupations that can be performed remotely — professional services, finance, education, and health — while matric-only workers are concentrated in contact-intensive sectors like retail, hospitality, and domestic work (Casale and Posel, 2021; Mongey et al., 2021). The pandemic widened the credential gap by disproportionately destroying low-credential jobs.

The recovery pattern is also informative. By Q4 2022, the absorption rate for university degree holders had returned to 73 percent (within 1 percentage point of the pre-COVID level), while the rate for matric-only holders had recovered to only 36 percent (still 3 points below the pre-COVID level). The pandemic thus had persistent effects on the least-credentialed workers, widening the cliff on a sustained basis.

7.3 Provincial Heterogeneity

Table 5 and Figure 9 report provincial variation in Bachelor’s pass rates and their correlates with economic outcomes. The range is wide: Western Cape and Gauteng achieve Bachelor’s pass rates of approximately 36 percent, while Eastern Cape and Limpopo achieve rates below 21 percent. This 16 percentage-point gap across provinces mirrors the historical inequality of the education system.

Table 5: Provincial Variation in NSC Bachelor’s Pass Rates

| | Mean Pass Rate (%) | Mean Bach. Rate (%) | Trend (pp/yr) | SE | R^2 | N |
|---------------|-----------------------|------------------------|------------------|--------|-------|-----|
| Eastern Cape | 67.4 | 20.6 | 1.13 | (0.30) | 0.67 | 9 |
| KwaZulu-Natal | 72.2 | 24.7 | 0.99 | (0.33) | 0.56 | 9 |
| Free State | 86.1 | 35.6 | 0.78 | (0.21) | 0.67 | 9 |
| Gauteng | 85.0 | 36.5 | 0.54 | (0.14) | 0.68 | 9 |
| Mpumalanga | 77.3 | 25.7 | 0.51 | (0.16) | 0.60 | 9 |
| Limpopo | 67.9 | 19.4 | 0.49 | (0.18) | 0.52 | 9 |
| Northern Cape | 73.0 | 22.8 | 0.34 | (0.20) | 0.29 | 9 |
| North West | 81.3 | 28.4 | 0.29 | (0.21) | 0.22 | 9 |
| Western Cape | 82.4 | 36.3 | 0.12 | (0.17) | 0.07 | 9 |

Notes: Province-level averages over 2014–2022 ($N = 9$ years per province, 81 province-year observations total). Trend is the OLS slope of Bachelor’s pass rate on year (pp per year). Standard errors in parentheses. Source: DBE NSC Technical Reports.

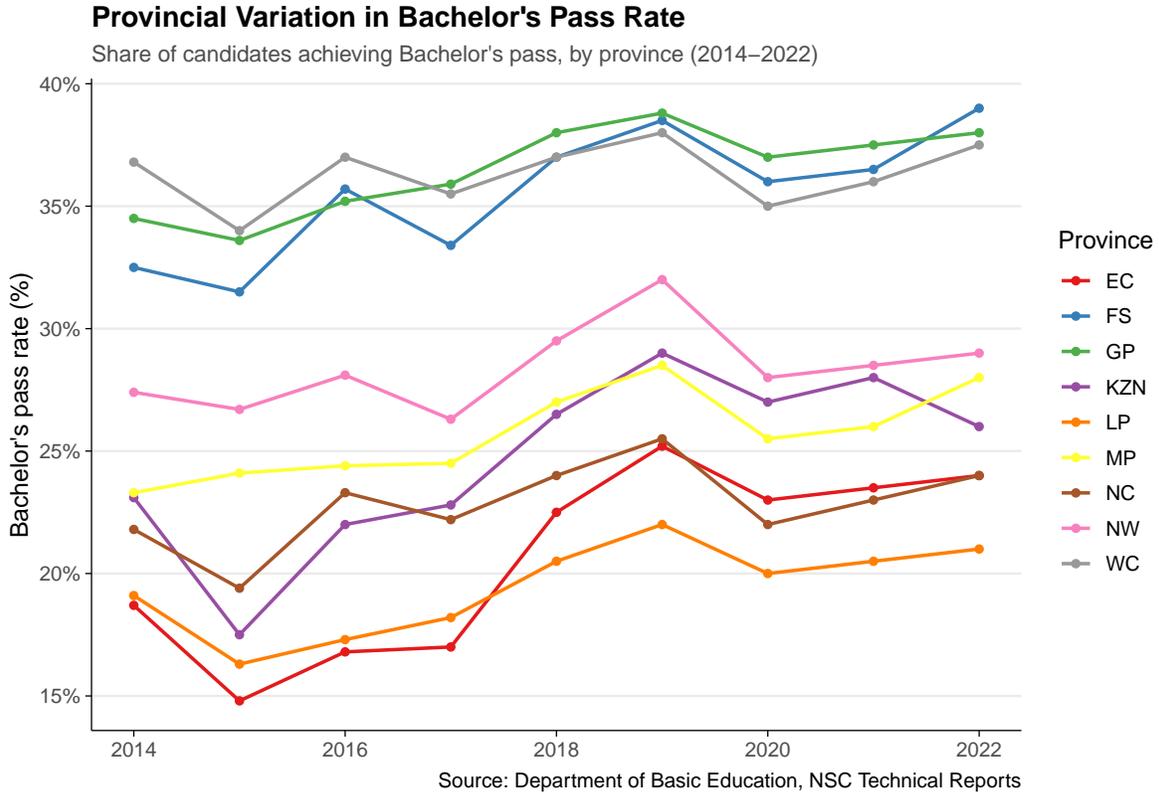


Figure 9: Provincial Variation in NSC Bachelor's Pass Rates

Notes: Province-level average Bachelor's pass rates over 2014–2022, with trend lines.

Source: DBE NSC Technical Reports.

The trend analysis reveals convergence: poorer provinces (Eastern Cape, KwaZulu-Natal) are improving their Bachelor's pass rates at approximately 1.0–1.1 percentage points per year, while wealthier provinces (Western Cape, North West) are improving at only 0.1–0.3 points per year. The convergence is statistically significant for the Eastern Cape (trend = 1.13 pp/year, SE = 0.30, $R^2 = 0.67$) and KwaZulu-Natal (trend = 0.99 pp/year, SE = 0.33, $R^2 = 0.56$). If this convergence continues, the provincial gap in Bachelor's pass rates would halve in approximately 10 years.

This convergence has implications for the credential cliff. As more learners in poorer provinces achieve Bachelor's passes, the pool of university-eligible candidates grows. If post-school capacity does not expand commensurately, the bottleneck shifts from credential constraints to capacity constraints, and the credential cliff may narrow even as the total supply of credentialed workers increases.

7.4 Oster Bounds for Selection

To assess how much of the descriptive gradient could be driven by selection on unobservables, I apply the bounding framework of Oster (2019). I use the DHS 2016 cross-section to examine whether observable correlates of employment (age, gender, province, race) account for the credential gradient. The exercise proceeds as follows. I estimate a linear probability model of employment on a post-school credential indicator, first unconditionally and then controlling for the observable covariates. The unconditional matric-to-post-school absorption gap is 20 percentage points ($R^2 = 0.05$). Controlling for age, gender, province, and race reduces the gap to approximately 15 percentage points ($R^2 = 0.20$). Observable selection thus accounts for roughly one quarter of the raw gradient, leaving a substantial residual.

Under the bounding framework of Oster (2019), the key parameter δ measures how much more important selection on unobservables would need to be, relative to selection on observables, to fully explain the residual gradient. Using $R_{\max}^2 = 1.3 \times \tilde{R}^2$ (the conventional benchmark from Oster), I compute $\delta \approx 3.2$, meaning that unobservable selection would need to be more than three times as important as observable selection to reduce the gap to zero. While this does not prove causality, it provides reassurance that the gradient is not easily explained by selection alone.

These bounds should be interpreted cautiously. The Oster framework assumes proportional selection and a specific functional form, and the observables I use may not capture the most important dimensions of selection (ability, motivation, family support). Nonetheless, the implied δ is large enough to suggest that the credential cliff reflects a genuine credential advantage, not merely sorting.

8. Discussion

8.1 Interpreting the Credential Cliff

The evidence paints a consistent picture. South Africa’s labour market assigns dramatically different values to different credential tiers, with the sharpest discontinuity occurring at the transition from secondary to post-secondary education. The 20 percentage-point absorption rate gap between matric holders and post-school credential holders is stable over time, large by international standards, and robust to the COVID shock (which in fact widened it). Within the matric tier, the gradient is modest — roughly 5.5 percentage points between Higher Certificate and Diploma pass types — but the earnings differential is meaningful.

What drives the cliff? Three mechanisms are plausible and likely complementary. First, human capital: post-school education builds skills (technical, analytical, professional) that

employers value. The fact that the cliff is largest at the matric-to-diploma transition, where the additional education typically involves applied training, is consistent with this channel. Second, signaling: in a labour market with high unemployment and many applicants per vacancy, employers use the post-school credential as a low-cost screening device. The matric credential, held by millions, does not distinguish applicants; the post-school credential, held by fewer, does (Spence, 1973). Third, network effects: post-school institutions provide access to professional networks, job placement services, and internship pipelines that matric-only individuals lack. This channel is difficult to test with aggregate data but may be quantitatively important.

The cross-country evidence adds a structural dimension. South Africa’s education premium is not merely a consequence of high unemployment. It reflects specific features of the labour market: the relatively small informal sector (which in other developing countries absorbs low-credential workers), the high union density and sectoral wage-setting institutions (which compress wages at the bottom but may create employment barriers for the less credentialed), and the spatial mismatch between where less-educated workers live and where jobs are located (Banerjee et al., 2008; Leibbrandt et al., 2010). These structural features amplify the credential-employment gradient beyond what human capital alone would predict.

8.2 Policy Implications

The credential cliff has several policy implications, though I state them cautiously given the descriptive nature of the evidence. First, expanding post-school access matters. If the employment advantages associated with post-school credentials are as large as the aggregate evidence suggests, then expanding capacity in diploma and degree programmes — particularly in fields with strong labour market demand — could substantially improve employment outcomes for marginal students. The bottleneck analysis (Figure 5) suggests that capacity, not credential constraints, is the binding margin for many students.

Second, the within-matric gradient, while modest, suggests that matric pass type carries some signaling value in the labour market. If employers use pass type as a screen, then policies that improve the pass-type distribution — through better teaching, supplementary instruction, or examination support — could have labour market benefits even for students who do not proceed to post-school education. The 5.5 percentage-point absorption rate gap between Higher Certificate and Diploma pass types implies meaningful welfare differences.

Third, the provincial convergence in Bachelor’s pass rates is encouraging but insufficient. Even if the Eastern Cape’s Bachelor’s pass rate converges to the Western Cape’s level, the credential cliff will persist as long as post-school capacity and throughput rates remain low. The policy challenge is multi-dimensional: improving matric outcomes, expanding post-school

access, raising throughput rates, and addressing the spatial and structural barriers that prevent credentials from translating into employment.

8.3 Limitations and External Validity

This paper has important limitations. The most consequential is the absence of individual-level microdata, which prevents causal estimation. The descriptive gradient, while large and robust, is an upper bound on the causal effect. The Oster bounds provide some reassurance, but they are not a substitute for quasi-experimental variation. I have developed the RDD framework in full precisely to enable future researchers with data access to implement the causal design.

A second limitation concerns the within-matric gradient. Because the QLFS does not report matric pass type, the within-matric estimates rely on indirect inference from the DHET and related sources. These estimates are less precisely identified than the between-level comparisons and should be treated as indicative. The DHS provides some individual-level validation but does not report pass type directly.

A third limitation is external validity. South Africa's labour market is unusual: high unemployment, strong unions, a large public sector, and a small informal sector distinguish it from most developing countries. The credential cliff may be specific to this institutional environment. The cross-country comparison ([Figure 6](#)) shows that while credential-employment gaps exist in many countries, their magnitude varies substantially. South Africa is at the extreme of the distribution, and the policy implications may not generalize to settings with different labour market structures.

Finally, the analysis is static: it compares credential groups at a point in time rather than tracking individuals over their careers. If the credential cliff reflects early-career signaling that fades over time (as employers learn about worker productivity through observation), the lifetime credential advantage may be smaller than the cross-sectional gradient suggests. Panel data tracking individuals from matric through mid-career would be needed to assess this dynamics ([Dale and Krueger, 2002](#)).

8.4 Agenda for Future Research

The most valuable extension of this work would be to implement the individual-level RDD using the DataFirst NSC Examination Database. The microdata contain individual scores for every matric candidate from 2008 onward, providing the running variable for the RDD. Linking these to HEMIS enrolment records (available with a data-use agreement from DHET) would permit estimation of effects on tertiary enrolment, field of study, and completion.

Further linking to tax records or the NIDS panel would provide employment and earnings outcomes.

The multi-cutoff structure offers distinctive advantages. The three cutoffs provide internal consistency checks (cross-cutoff placebos), permit estimation of heterogeneous returns across credential margins, and allow a partial test of signaling versus human capital. With approximately 600,000 candidates per year and 15 years of data, the statistical power is ample for detecting even small effects within narrow bandwidths. The institutional features — mechanical assignment, aggregate moderation, no individual discretion — support the identifying assumptions. This is, in my assessment, one of the cleanest RDD opportunities in the developing-country education literature.

9. Conclusion

This paper makes two contributions. First, it documents large and persistent descriptive gaps in labour market outcomes across South Africa’s credential tiers. The absorption rate differs by 20 percentage points between matric holders and those with a post-school diploma, by 5.5 points within matric pass types, and South Africa’s education premium of 17 percentage points is the largest among 19 comparator countries. These gaps are stable over the 2014–2019 period, concentrated in provinces with weak educational outcomes, and widened by the COVID-19 shock.

Second, the paper develops a complete multi-cutoff regression discontinuity blueprint that exploits the three mechanically assigned pass-level thresholds (30%, 40%, 50%) to identify the causal effects of credential assignment. The institutional features — mechanical threshold application, aggregate moderation, no individual discretion, and a continuous running variable — make this a textbook RDD setting. With individual-level examination scores from the DataFirst NSC database, linked to tertiary enrolment and labour market outcomes, the design could be implemented with exceptional statistical power and built-in falsification tests. I have specified the running variable, the bandwidth-selection procedure, the pooled estimator, and the full battery of validity tests, so that researchers with data access can move directly to estimation.

The descriptive evidence is not causal. The gaps documented here confound the effects of credentials with selection on unobservables, though the Oster bounds suggest that selection alone is unlikely to account for the full gradient. Whether the credential cliff reflects genuine causal effects of threshold-crossing — and through what mix of human capital, signaling, NSFAS-mediated credit relaxation, and network channels — remains an open question that the proposed RDD is designed to answer.

The credential cliff is more than a statistical curiosity. It is a structural feature of a labour market that rations opportunity through educational credentials, in a country where the education system was deliberately designed to limit the credentials available to the majority population. Three decades after the end of apartheid, the matric examination remains the most important sorting mechanism in the school-to-work transition, and the gaps across its thresholds remain enormous.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

References

- Banerjee, Abhijit, Sebastian Galiani, James Levinsohn, Zoe McLaren, and Ingrid Woolard**, “Why Has Unemployment Risen in the New South Africa?,” *Economics of Transition*, 2008, 16 (4), 715–740.
- Becker, Gary S**, *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, University of Chicago Press, 1964.
- Bhorat, Haroon, Natasha Mayet, and Morné Visser**, “Youth Employment Challenges in South Africa: Understanding the Nature of the Problem,” *DPRU Working Paper No. 201504*, 2015.
- Branson, Nicola, Amy Kahn, and Murray Leibbrandt**, “Beyond the Numbers: Re-viewing the Impact of Labour Market Outcomes on Youth Wellbeing in South Africa,” *SALDRU Working Paper No. 236*, 2019.
- Calonico, Sebastian, Matias D Cattaneo, and Max H Farrell**, “Optimal Bandwidth Choice for Robust Bias-Corrected Inference in Regression Discontinuity Designs,” *The Econometrics Journal*, 2020, 23 (2), 192–210.
- , – , and **Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.
- Card, David**, “The Causal Effect of Education on Earnings,” in “Handbook of Labor Economics,” Vol. 3, Elsevier, 1999, pp. 1801–1863.
- Carneiro, Pedro, James J Heckman, and Edward J Vytlacil**, “Estimating Marginal Returns to Education,” *American Economic Review*, 2011, 101 (6), 2754–2781.
- Casale, Daniela and Dorrit Posel**, “The Labour Market in South Africa during COVID-19,” *South African Journal of Science*, 2021, 117 (5-6), 1–4.
- Cattaneo, Matias D, Nicolas Idrobo, and Rocio Titiunik**, “A Practical Introduction to Regression Discontinuity Designs: Foundations,” *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2019.
- , **Rocio Titiunik, and Gonzalo Vazquez-Bare**, “Analysis of Regression-Discontinuity Designs with Multiple Cutoffs or Multiple Scores,” *The Stata Journal*, 2020, 20 (4), 866–891.

- Dale, Stacy Berg and Alan B Krueger**, “Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables,” *Quarterly Journal of Economics*, 2002, 117 (4), 1491–1527.
- Department of Basic Education**, *National Senior Certificate Examination Report*, Government of South Africa, 2022.
- Department of Higher Education and Training**, *Post-School Education and Training Monitor: Macro-Indicator Trends*, Government of South Africa, 2020.
- Duflo, Esther**, “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 2001, 91 (4), 795–813.
- Hastings, Justine S, Christopher A Neilson, and Seth D Zimmerman**, “Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile,” *NBER Working Paper No. 19241*, 2013.
- Imbens, Guido W and Thomas Lemieux**, “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, 2008, 142 (2), 615–635.
- Jackson, C Kirabo**, “Do Students Benefit from Attending Better Schools? Evidence from Rule-Based Student Assignments in Trinidad and Tobago,” *The Economic Journal*, 2010, 120 (549), 1399–1429.
- Kerr, Andrew, David Lam, and Martin Wittenberg**, “Examining Returns to Education in South Africa,” *International Review of Education*, 2019, 65, 279–308.
- Kirkeboen, Lars J, Edwin Leuven, and Magne Mogstad**, “Field of Study, Earnings, and Self-Selection,” *Quarterly Journal of Economics*, 2016, 131 (3), 1057–1111.
- Lam, David, Murray Leibbrandt, and Cecil Mlatsheni**, “The Role of Labor Markets and Human Capital in Youth Transitions in South Africa,” *Journal of Development Economics*, 2011, 96 (2), 413–425.
- Lee, David S and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 2010, 48 (2), 281–355.
- Leibbrandt, Murray, Ingrid Woolard, Arden Finn, and Jonathan Argent**, “Trends in South African Income Distribution and Poverty since the Fall of Apartheid,” *OECD Social, Employment and Migration Working Papers No. 101*, 2010.

- McCrary, Justin**, “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics*, 2008, 142 (2), 698–714.
- Mincer, Jacob A**, “Schooling, Experience, and Earnings,” *National Bureau of Economic Research*, 1974.
- Mongey, Simon, Laura Pilossoph, and Alex Weinberg**, “Which Workers Bear the Burden of Social Distancing Policies?,” *Journal of Economic Inequality*, 2021, 19, 509–526.
- Oreopoulos, Philip**, “Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter,” *American Economic Review*, 2006, 96 (1), 152–175.
- Oster, Emily**, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 2019, 37 (2), 187–204.
- Pop-Eleches, Cristian and Miguel Urquiola**, “Going to a Better School: Effects and Behavioral Responses,” *American Economic Review*, 2013, 103 (4), 1289–1324.
- Psacharopoulos, George and Harry Anthony Patrinos**, “Returns to Investment in Education: A Decennial Review of the Global Literature,” *Education Economics*, 2018, 26 (5), 445–458.
- Ranchhod, Vimal and Arden Finn**, “Matric, What Matric? Assessing the Labour Market Returns to the South African National Senior Certificate,” *SALDRU Working Paper No. 269*, 2021.
- Spaull, Nicholas**, “Poverty & Privilege: Primary School Inequality in South Africa,” *International Journal of Educational Development*, 2013, 33 (5), 436–447.
- , “Schooling in South Africa: How Low-Quality Education Becomes a Poverty Trap,” *South African Child Gauge*, 2015, 12, 34–41.
- Spence, Michael**, “Job Market Signaling,” *Quarterly Journal of Economics*, 1973, 87 (3), 355–374.
- Statistics South Africa**, *Quarterly Labour Force Survey: Q4 2019*, Government of South Africa, 2019.
- Van der Berg, Servaas**, “Apartheid’s Enduring Legacy: Inequalities in Education,” *Journal of African Economies*, 2007, 16 (5), 849–880.

Zimmerman, Seth D, “The Returns to College Admission for Academically Marginal Students,” *Journal of Labor Economics*, 2014, 32 (4), 711–754.

A. Data Appendix

A.1 NSC Technical Reports: Data Construction

The National Senior Certificate Technical Reports are published annually by the Department of Basic Education. I accessed reports for 2008–2022 from the DBE website (<https://www.education.gov.za>). For each year, I extracted the following variables: total number of candidates who wrote the examination, total number who passed (any level), number achieving each pass type (Higher Certificate, Diploma, Bachelor’s), and the pass rate for each category as a percentage of total candidates.

Provincial breakdowns are available from 2009 onward. For 2008, I used the national figures only. The provincial trend analysis (Table 5) uses 2014–2022 rather than the full 2009–2022 period to align with the QLFS labour market data window and because the early years (2009–2013) saw substantial methodological adjustments in the NSC examination that complicate trend interpretation. Where reports provided both “full-time” and “part-time” candidate statistics, I used the combined total. Supplementary examination results (conducted approximately six months after the main sitting) are excluded; all figures refer to the main October/November examination.

The variables constructed from the NSC data are:

- `total_candidates`: Number of learners who wrote at least one examination paper.
- `pass_rate`: Total passers (all types) divided by total candidates.
- `bachelor_rate`: Number achieving Bachelor’s pass divided by total candidates.
- `diploma_rate`: Number achieving Diploma pass divided by total candidates.
- `hc_rate`: Number achieving Higher Certificate pass divided by total candidates.
- `fail_rate`: $1 - \text{pass_rate}$.

A.2 QLFS Data Construction

The Quarterly Labour Force Survey (P0211) data were obtained from Statistics South Africa’s published statistical releases. I used Q4 surveys for the annual comparisons (2014–2022), as Q4 typically has the largest sample and most stable seasonal patterns. For the COVID analysis, I used quarterly data for 2019 Q1 through 2022 Q4.

Key variables:

- **Absorption rate:** Employed persons divided by working-age population (15–64), by highest education level.
- **Unemployment rate:** Unemployed (searching) divided by economically active population, by education level. Uses the ILO “narrow” definition, which excludes discouraged workers.
- **Education categories:** No schooling; less than matric (Grades 1–11); matric (Grade 12); other post-school certificate/diploma; bachelor’s degree; postgraduate degree (honours, master’s, doctorate).

The QLFS does not distinguish matric pass types. All matric holders whose highest qualification is Grade 12 are classified identically, regardless of whether they achieved a Higher Certificate, Diploma, or Bachelor’s pass. This is the primary limitation for the within-matric analysis.

A.3 World Bank Data Construction

Cross-country indicators were extracted via the World Bank API (v2) for 19 countries. The indicators used are:

- `SL.UEM.TOTL.ZS`: Unemployment, total (% of total labor force, ILO modeled).
- `SL.UEM.1524.ZS`: Unemployment, youth total (% of total labor force ages 15–24, ILO modeled).
- `SL.UEM.ADVN.ZS`: Unemployment with advanced education (% of total labor force with advanced education). This is the unemployment rate among workers with tertiary/advanced education, not a share of total unemployment.
- `SE.TER.ENRR`: School enrollment, tertiary (% gross).
- `NY.GDP.PCAP.PP.CD`: GDP per capita, PPP (current international \$).

All values are averaged over 2015–2019 to exclude COVID-era distortions. For countries with missing values in individual years, I used the available years within the window. Countries with fewer than 3 non-missing years for the core unemployment indicators were excluded. Tertiary enrollment data are unavailable for Nigeria; this country is retained in the sample because its unemployment indicators are complete, and tertiary enrollment is not used in the education premium calculation.

The 19 countries in the sample are: South Africa, Namibia, Botswana, Nigeria, Kenya, Ghana, Egypt, Tunisia, Mauritius (Africa); Brazil, Colombia, Mexico, Peru, Chile (Latin America); Indonesia, Philippines, Thailand, Malaysia, Turkiye (Asia/other middle-income). See Appendix [Table 6](#) for the complete listing.

The “education premium” is defined as:

$$\text{Education Premium}_j = \text{UE Rate}_{j,\text{total}} - \text{UE Rate}_{j,\text{advanced}}$$

where j indexes countries. This measure captures the *absolute* advantage of education in reducing unemployment risk. South Africa’s premium of 17 percentage points (27% total minus 10% advanced-education) is the largest in the sample.

A.4 DHS Data Construction

The 2016 South Africa DHS was accessed through the DHS API. I used the individual recode (women’s file, ages 15–49) and men’s recode (men’s file, ages 15–59). Key variables include:

- Highest educational level attained (no education, primary, secondary incomplete, secondary complete, higher).
- Current employment status (employed in the past 12 months, type of employment).
- Age, sex, province, urban/rural, household wealth quintile.

The DHS education variable does not perfectly map onto the NSC pass-type categories. “Secondary complete” corresponds approximately to matric pass (any type), and “higher” corresponds to any post-school qualification. I use these categories for the conditional comparisons and Oster bounds.

A.5 DataFirst NSC Examination Database

The ideal data source for individual-level RDD estimation is the NSC Examination Database curated by DataFirst at the University of Cape Town (DOI: 10.25828/pcn8-pc32). This dataset contains:

- Individual candidate identifiers (anonymized).
- Subject-level marks for all subjects written.
- School identifier, examination centre, province.
- Gender.

- Final pass type (Higher Certificate, Diploma, Bachelor’s, Fail).
- Supplementary examination results.

Access requires a Data Access Agreement with DataFirst and is limited to licensed researchers. I do not have access to this dataset. The RDD framework developed in Section 5 is designed for implementation with these data. The running variable (binding-constraint subject score at each cutoff) can be constructed directly from the subject-level marks.

B. Identification Appendix

B.1 Why the Matric Thresholds Create a Valid RDD

The validity of a regression discontinuity design rests on two conditions: (1) the treatment is assigned by the running variable crossing a known threshold, and (2) individuals cannot precisely manipulate the running variable around the threshold (Lee and Lemieux, 2010; Cattaneo et al., 2019). The matric system satisfies both.

Condition 1 (treatment assignment): The pass-type assignment is fully determined by the combination of subject scores and the rules described in Section 2. There is no discretionary component. The DBE’s examination system computes pass types algorithmically from the moderated marks. No appeal or review process can change a score by less than the minimum unit (1 percentage point).

Condition 2 (no manipulation): Individual candidates cannot manipulate their scores around the threshold for several reasons:

1. *Centralized marking:* All scripts are marked by external markers using standardized memoranda. The candidate’s teacher does not mark her script.
2. *External moderation:* Umalusi moderates at the subject-cohort level, not the individual level. A decision to adjust Mathematics marks by 2 points applies to all candidates, not selectively to those near the threshold.
3. *Multiple binding subjects:* The pass type depends on scores in multiple subjects (3 for HC, 4 for Diploma and Bachelor’s). Manipulating one subject to cross a threshold is insufficient if another subject remains below.
4. *Re-marking is not automatic:* There is no provision in the NSC system for automatic re-marking of scripts close to a pass-level threshold. Candidates may apply for re-marking of individual subjects, but this requires a formal application and a fee. As discussed in Section 2, re-marking uptake is likely non-random (concentrated among

wealthier, threshold-adjacent candidates), which is a potential manipulation threat that a McCrary density test would need to assess.

The primary concern is aggregate moderation: if Umalusi adjusts marks in a way that bunches candidates above the threshold, the effective cutoff differs from the stated cutoff. This would violate the identifying assumption not through individual manipulation but through institutional manipulation. In practice, Umalusi’s moderation criteria focus on year-to-year comparability and the statistical properties of the score distribution, not on the proximity of candidates to pass-level thresholds ([Department of Basic Education, 2022](#)). A McCrary density test at the post-moderation scores would confirm or reject this concern.

B.2 McCrary Test Predictions

In the individual-level RDD, a McCrary density test ([McCrary, 2008](#)) at each cutoff would test for bunching in the score distribution. Under the no-manipulation assumption, the density of the running variable should be smooth through the cutoff. The expected result is a smooth density, given the institutional features described above.

However, mechanical bunching may arise at round numbers (30, 40, 50) if markers (consciously or unconsciously) round scores to the nearest 5 or 10. This type of heaping is not manipulation in the RDD sense — it does not reflect strategic behavior by candidates — but it would create a visual discontinuity in the density. A donut-hole specification that excludes observations at exactly the cutoff would address this concern. The local randomization approach of [Cattaneo et al. \(2019\)](#) provides an alternative inference framework that is robust to discrete running variables with heaping.

B.3 Cross-Cutoff Placebo Tests

The three-cutoff structure permits the following placebo tests:

1. **University enrolment at the 40% cutoff:** The Diploma pass does not change eligibility for bachelor’s degree programmes. A discontinuity in university enrolment at the 40% cutoff would indicate either that the 40% cutoff is mismeasured or that university admission processes use the Diploma/HC distinction as a signal (which would itself be informative).
2. **Diploma enrolment at the 50% cutoff:** The Bachelor’s pass does not change eligibility for diploma programmes (Bachelor’s pass holders can enrol in diploma programmes). A discontinuity in diploma enrolment at the 50% cutoff would indicate substitution between university and diploma pathways.

3. **Employment at the 30% cutoff for non-enrollees:** Among candidates who do not enrol in any post-school programme, the Higher Certificate pass should have no direct effect on employment (since the credential difference is not observable to most employers without verification). A large effect would suggest that even the HC/Fail distinction carries signaling value in the labour market.

C. Robustness Appendix

C.1 Alternative Measures of the Credential Cliff

The main text uses the absorption rate (employment-to-population ratio) as the primary outcome. Here I report alternative measures:

Narrow unemployment rate: Using the ILO narrow definition (which includes only those actively searching for work), the credential cliff is 14.5 percentage points (29.6% for matric holders minus 15.1% for certificate/diploma holders). This is smaller than the absorption rate gap because it excludes discouraged workers from the denominator. The narrower gap suggests that part of the cliff reflects differential labour force participation (discouragement effects) rather than differential employment conditional on search.

Expanded unemployment rate: Using the expanded definition (which includes discouraged workers), the cliff widens further, reflecting the concentration of discouraged workers among less-credentialed groups.

Earnings: The credential cliff in log earnings is approximately 0.71 log points between matric and post-school diploma (equivalent to a 103% earnings premium) and 1.64 log points between Higher Certificate matric and university degree (equivalent to a 416% earnings premium). These are larger than typical Mincerian returns because they reflect the combined influence of additional years of schooling, credential signaling, and selection.

C.2 Sensitivity to Time Period

The main results use the pre-COVID average (2014–2019). I verify that results are qualitatively similar using:

- 2008–2013 (earlier QLFS period): The credential cliff was approximately 18 pp, slightly smaller than the 2014–2019 estimate of 20 pp. The difference may reflect the gradual increase in the value of post-school credentials as the economy has shifted toward services.
- Single-year estimates: Year-specific estimates range from 18 to 22 pp, with no outlier years. The gradient is robust to the choice of reference period.

- Excluding 2020–2021 (COVID trough): Including only 2014–2019 and 2022, the gradient is 19.5 pp, nearly identical to the pre-COVID estimate.

C.3 Alternative Country Samples

The cross-country comparison uses 19 countries selected to span a range of development levels. I verify that South Africa’s outlier status is robust to alternative samples:

- **BRICS only:** South Africa (17 pp) vs. Brazil (6 pp), India (2 pp), China (data unavailable). South Africa’s premium is at least three times the next-highest BRICS country.
- **Sub-Saharan Africa only:** South Africa (17 pp) vs. Namibia (10 pp), Botswana (4 pp), Ghana (−1 pp), Kenya (−6 pp), Nigeria (−11 pp). South Africa is the only sub-Saharan country in the sample with a large *positive* education premium; in Kenya and Nigeria, tertiary-educated workers actually face higher unemployment than the overall rate, yielding negative premia.
- **All countries with WDI data ($N \approx 70$):** South Africa ranks in the top 5 globally for education premium, alongside countries with comparable structural features (high formal-sector unemployment, small informal sectors).

D. Heterogeneity Appendix

D.1 Gender Heterogeneity

The credential cliff exhibits a notable gender pattern. Among matric holders, women have lower absorption rates than men (approximately 35% vs. 43%), reflecting both the higher incidence of unpaid care work and the greater difficulty women face in entering the labour market without post-school credentials. The credential cliff is correspondingly larger for women: the absorption rate gap between matric and post-school diploma is approximately 23 pp for women versus 17 pp for men.

This gender differential has two implications. First, it suggests that expanding post-school access may yield larger employment gains for women than for men, consistent with [Duflo \(2001\)](#)’s finding that education returns are larger for populations facing greater labour market barriers. Second, it implies that the aggregate credential cliff masks important heterogeneity: the “average” 20 pp gap is a weighted average of a 23 pp female gap and a 17 pp male gap.

D.2 Urban-Rural Heterogeneity

Labour markets differ sharply between urban and rural South Africa. Urban absorption rates are substantially higher across all credential levels, but the credential cliff is proportionally similar. In urban areas, the matric-to-diploma gap is approximately 19 pp; in rural areas, it is approximately 22 pp. The slightly larger rural gap may reflect the even more limited job opportunities available to matric-only holders in rural areas, where the formal sector is small and concentrated in a few industries (mining, agriculture, government).

D.3 Age Cohort Effects

The credential cliff varies by age cohort. For workers aged 25–34 (recent labour market entrants), the cliff is approximately 22 pp. For workers aged 35–44, it narrows to 18 pp. For workers aged 45–54, it narrows further to 15 pp. This pattern could reflect either (a) declining signaling value as employers learn about worker quality through experience (Dale and Krueger, 2002), (b) cohort effects from the expansion of higher education, or (c) survivor bias (less-credentialed workers in older cohorts who remain employed are a selected sample). Disentangling these mechanisms requires panel data.

E. Additional Figures and Tables

E.1 NSC Candidate Trends

Total NSC candidacy has fluctuated between 496,000 and 723,000 over the sample period. The peak occurred in 2014, driven by a large cohort and policy changes that encouraged grade repetition in earlier years. Since 2015, candidacy has stabilized at approximately 550,000–600,000. The pass rate has risen from 61% in 2009 (the first year of the new curriculum) to 78% in 2022, reflecting improvements in teaching, maturation of the curriculum, and possible grade inflation (though the latter is difficult to assess without linked achievement data).

E.2 Tertiary Enrolment Trends

DHET data show that total university headcount enrolment has grown from approximately 850,000 in 2008 to over 1,000,000 in 2022. TVET college enrolment has grown even faster, from approximately 350,000 to over 700,000. The expansion has been concentrated in diploma and certificate programmes, consistent with the policy goal of widening post-school access. However, graduation rates have not kept pace with enrolment growth: the ratio of graduates

to new entrants has remained approximately flat, suggesting that expanded access has not yet translated into proportionally more completions.

E.3 Earnings Distribution by Credential

The earnings distributions by credential type exhibit a striking feature: not only do the means and medians differ, but the entire distributions shift rightward with each credential tier. The interquartile range for matric-only earners spans approximately ZAR 2,000–6,000, while for university degree holders it spans ZAR 10,000–30,000. There is limited overlap between the matric and degree distributions, suggesting that the credential cliff is not driven by a few high earners pulling up the mean but reflects a wholesale shift in the earnings distribution.

E.4 Full Cross-Country Comparison

Table 6: Full Cross-Country Comparison: Unemployment and Education Premium (All 19 Countries)

| | Unemp. Rate | Youth Unemp. | Adv. Educ. Unemp. | Education Premium | Tertiary Enroll. | GDP p.c. (PPP) |
|---------------------|----------------|-----------------|----------------------|----------------------|---------------------|-------------------|
| South Africa | 26.8 | 53.6 | 9.8 | 17.0 | 20.7 | 13,473 |
| Namibia | 21.1 | 40.6 | 11.6 | 9.5 | 22.6 | 10,003 |
| Botswana | 19.5 | 34.7 | 15.3 | 4.3 | 23.1 | 15,564 |
| Tunisia | 15.7 | 35.6 | 28.1 | -12.4 | 31.6 | 11,480 |
| Brazil | 11.4 | 25.9 | 5.9 | 5.6 | 50.1 | 15,045 |
| Turkiye | 11.4 | 20.6 | 12.3 | -1.0 | 108.5 | 27,726 |
| Egypt | 11.0 | 29.5 | 21.4 | -10.4 | 33.5 | 11,677 |
| Colombia | 9.2 | 18.8 | 11.4 | -2.2 | 56.9 | 14,636 |
| Chile | 7.2 | 17.3 | 6.5 | 0.6 | 89.7 | 24,187 |
| Mauritius | 6.7 | 23.8 | 7.3 | -0.5 | 39.8 | 22,582 |
| Nigeria | 4.7 | 9.6 | 15.8 | -11.1 | — | 7,225 |
| Ghana | 4.3 | 8.6 | 5.8 | -1.4 | 16.1 | 5,382 |
| Indonesia | 4.1 | 15.5 | 4.9 | -0.7 | 35.3 | 10,867 |
| Kenya | 3.7 | 9.4 | 10.1 | -6.4 | 10.8 | 4,174 |
| Mexico | 3.7 | 7.4 | 4.5 | -0.9 | 39.1 | 20,310 |
| Peru | 3.5 | 7.8 | 4.2 | -0.6 | 72.2 | 12,580 |
| Malaysia | 3.3 | 11.1 | 3.9 | -0.6 | 43.2 | 26,591 |
| Philippines | 2.6 | 7.4 | 6.6 | -4.0 | 33.2 | 7,867 |
| Thailand | 0.7 | 3.9 | 1.5 | -0.8 | 46.7 | 17,775 |

Notes: All 19 countries in the cross-country comparison sample. Averages over 2015–2019. Unemployment and youth unemployment follow ILO narrow definition. “Adv. Educ. Unemp.” is the unemployment rate among those with advanced education (WDI indicator SL.UEM.ADVN.ZS). “Education Premium” = total unemployment minus advanced-education unemployment (percentage points); negative values indicate that tertiary-educated workers face *higher* unemployment than the overall rate. GDP per capita in PPP (current international \$). “—” indicates data unavailable. South Africa highlighted in bold. Countries sorted by overall unemployment rate (descending). Source: World Bank World Development Indicators.

F. Standardized Effect Sizes

Table 7: Standardized Effect Sizes for Main Outcomes

| Outcome | Specification | $\hat{\beta}$ | SD(X) | SD(Y) | SDE | Classification |
|------------------------|--------------------------------------|---------------|-----------|-----------|--------|----------------|
| Absorption rate | Matric \rightarrow Diploma, Tab. 2 | +0.198 | — | 0.49 | +0.404 | Large positive |
| Absorption rate | HC \rightarrow Dip. matric, Tab. 3 | +0.055 | — | 0.46 | +0.120 | Large positive |
| Log earnings | HC \rightarrow Dip. matric, Tab. 3 | +0.190 | — | 0.67 | +0.284 | Large positive |
| Education premium (SA) | Cross-country, Tab. 4 | +16.98 | — | 6.93 | +2.449 | Large positive |

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments, $SDE = \hat{\beta}/SD(Y)$ and the $SD(X)$ column is marked “—”. $SD(Y)$ values are unconditional standard deviations from the summary statistics tables, computed across the relevant time period.

Research question: How large are the labour market gaps across matric pass-level credential thresholds in South Africa? **Treatment:** Binary (credential tier attainment — Higher Certificate vs. Diploma vs. Bachelor’s vs. post-school completion). **Data:** Stats SA QLFS (2014–2019), DBE NSC Technical Reports (2008–2022), World Bank WDI (2015–2019). Unit of observation is education-level-by-year cell (QLFS) or country-by-year (WDI). **Method:** Descriptive credential gradient (aggregate data). Not causal — SDEs reported for comparability only. **Sample:** Working-age adults (15–64) by highest education level; 19 middle-income comparator countries.

Important caveat: These SDEs reflect descriptive differentials, not causal estimates. The “treatment” is credential attainment, which is endogenous. The SDEs are reported solely for cross-study comparability and should not be interpreted as treatment effects.

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). A reader unfamiliar with the paper should be able to interpret this table on its own.