The Educational and Labour Market Effects of the 1976 Universal Primary Education Programme in Nigeria[†]

By Neron Sifflore[*]

In 1976, Nigeria introduced its free education programme that abolished fees in all public primary schools. Using data collected 33 years after the scheme's introduction, I estimate its impacts on educational attainment, learning, and employability. I find the programme increased men's educational attainment by one year but had no measurable effect on their learning; it also increased the probability of a woman completing primary school by 36.8 percent. I find no evidence that the programme improved employability. If anything, It seems to have reduced male participation in agriculture, but I find no evidence of increased participation in skilled jobs.

Over the past decades, many developing countries have instituted policies that abolish tuition fees or build more schools, partly to rejuvenate their education systems and to promote equitable access to education. Malawi, for example, abolished tuition fees in all public primary schools in 1994; Uganda did so in 1997, and Kenya, in 2003. Did these policies improve the learning outcomes of their beneficiaries? If so, did the increase in human capital translate into better labour market outcomes later in life?

An extensive literature addresses the first question—an important one because most sub-Saharan African countries that still levy fees in primary schools are considering abolishing them in the coming years. However, existing research on the effects of these policies on

[†] Nottingham School of Economics, University of Nottingham, Malaysia Campus, Jalan Broga, Semenyih, 43500 Selangor, Malaysia

[*] Email: neronsifflore@gmail.com

learning outcomes mainly focuses on one gender, commonly women. Focusing on one gender may overestimate or underestimate the effects of these policies if men and women respond differently to reductions in the costs of education. Moreover, although a few studies suggest that the beneficiaries of these policies may enjoy better employment outcomes once they enter the labour market, the literature does not say much about whether these effects persist in the long run.

In this paper, I examine the schooling and labour market effects of the 1976 Universal Primary Education (UPE) programme in Nigeria, which eliminated fees in all public primary schools and built over 12000 primary schools during the 1976 academic year. The programme coincided with a surge in primary school enrolment. Specifically, the number of children enrolled in primary schools increased from 4.4 million in 1974 to 13.8 million in 1981.

I identify the effects of the programme on men and women using a difference-in-differences strategy that leverages cross-cohort and regional variation in exposure to the programme. Cross-cohort variation comes from the fact that mainly children who were of primary school age during the programme's implementation could benefit from it, and regional variation stems from differences in the intensity of the programme across states. I proxy programme intensity using the gap between a state's potential and actual primary school enrolment in the pre-programme period. Intuitively, the programme could have little to no effect in a state where all children were already enrolled in primary school before its introduction, and the largest effects in states with the lowest pre-programme primary school enrolment rates.

I use data from a large nationally representative household survey, the 2009 Harmonised Nigeria Living Standards Survey. The survey was conducted 33 years after the programme was launched, thus facilitating the study of the effects on long-run labour market outcomes. I have roughly 20,000 men and 20,000 women in my sample with an average level of 6 and 4 years of education, respectively. Most men worked in agriculture, and around two in five held wage jobs. Around one in three women worked in agriculture, and one in ten held wage jobs.

I find the programme increased men’s educational attainment but had limited effects on their learning: it increased educational attainment by one year, a large effect, since the average educational attainment for men is around 6 years. The programme also increased the probability a man completed primary school, junior secondary school, and senior secondary school by 13.7 percent, 21.7 percent, and 22.3 percent, respectively. However, I find no evidence that the programme increased male enrolment in tertiary education. The programme also does not seem to have improved male employability: It reduced male participation in agriculture by around 11 percent, but I find no evidence of increased participation in skilled work, the public sector, or wage work.

Most of the effects on women are imprecisely estimated. Nevertheless, I find the programme increased the likelihood a woman completed primary school by 36.8 percent and the probability a woman can read in English by 28.6 percent. The latter estimate is statistically significant only at the ten percent level, however. The programme does not seem to have increased the probability a woman completed junior secondary or senior secondary school; I also find no evidence that it improved women’s employability.

I also use a cohort analysis that isolates the programme’s effect on each birth cohort. Consistent with the main difference-in-differences results, I find the programme improved the educational attainment of men born between 1970 and 1975, who were its main beneficiaries. The programme also seems to have improved the educational attainment of some male cohorts that were past primary school age in 1976, which is most likely due to overage enrolment.

In contrast to the main difference-in-differences results, however, the cohort analysis indicates that the programme had negative effects on the primary school completion rates of female cohorts that were of primary school age during its implementation. Motivated by evidence that Muslim families were against 'Western' education as a means of educating their girls at the time when the programme was launched, I explore the extent to which cultural factors may drive these results by estimating the programme's effects on Muslim and non-Muslim women separately. The negative effects seem to be concentrated among the former, which suggests that cultural prohibition against girls' education might have undermined the accumulation of human capital among Muslim girls.

I perform several robustness checks: I address selective migration; I test the plausibility of the key identifying assumption, the parallel trends assumption; I also run a series of placebo tests whereby I assume the programme was implemented in the years after its actual implementation. The main findings are robust to accounting for selective migration. Moreover, the parallel trends assumption seems to hold for most outcomes. However, I also find statistically significant effects from some of the placebo programmes, which suggests that I cannot rule out the possibility that unobserved factors could be driving some of my results.

This paper contributes to the literature in three ways. First, I examine how a programme that expands access to primary education affects the schooling and employment outcomes of men and women, which complements those studies focusing on only one gender. Second, I consider labour market outcomes over the long run, which complements research examining the short-run effects of such policies. Finally, I also provide novel evidence that the benefits of the programme were disproportionately enjoyed by boys and non-Muslim girls.

The remainder of this paper proceeds as follows. Section II reviews the literature. Section III describes the UPE programme. Section IV describes the identification strategy; section V the data; and Section VI presents estimates of the programme's effects on schooling and employment outcomes. Section VII tests the robustness of the paper's main findings. Lastly, section VIII concludes.

I. Literature Review

In this section, I review the relevant empirical literature on the educational and labour market effects of policies that expand access to education in developing countries. To gauge how these initiatives might affect educational and employment outcomes, I first present a brief conceptual framework.

A. Conceptual Framework

The conceptual framework I use in this paper draws on a model of parental investment in children's education developed by Orazem and King (2007). The theoretical foundations for this model trace back to early work on human capital theory (Mincer, 1958; Becker, 1963), which frames education as a long-term investment, with individuals weighing the costs and expected benefits when choosing the amount of schooling to pursue. Orazem and King's model assumes parents are utility-maximizing agents whose welfare depends on their household's consumption of goods and their children's human capital production. Parents decide how to allocate their sons' and daughters' time between work (e.g., household chores) to produce consumption goods and school to accumulate human capital.

The model incorporates insights from Becker's (1971) model of taste discrimination, allowing parents to value the education of their sons more than that of their daughters, even if the two have identical abilities and face identical school supply. This aspect is consistent with the reality in Nigeria, where cultural norms and prohibitions against educating girls exist (Csapo, 1983). The model predicts that an initiative that decreases the direct costs of schooling, like the UPE programme, will reduce the price of children's time spent in school relative to the price of consumption, inducing parents to send both their sons and daughters to school. However, the presence of cultural prohibitions against educating girls might mute the positive effect on the latter.

Extending these predictions to the labour market suggests that by raising the productivity of boys and girls, the initiative will increase their employability once they enter the labour force. However, several factors that scholars have established to undermine the efficiency of labour markets in developing countries—such as a lack of employment opportunities (Csapo, 1983; Duflo *et al.,* 2021), a scarcity of capital (Khatkhate, 1978), and information friction (Bassi and Nansamba, 2017)—will weaken this positive effect. The overall impact on employability will therefore depend on which of these two sets of opposing forces dominates.

B. Empirical Evidence

Studies on the effects of schooling expansions in developing countries find positive educational effects consistent with the above predictions. The impacts on labour market outcomes, on the other hand, are less documented and more mixed. These studies seek to understand how improving the accessibility of education affects several outcomes, including educational attainment (Duflo, 2001; Chicoine, 2012; Khanna, 2020; Chicoine, 2021; Brudevold-Newman, 2021; Nandi *et al.,* 2023), employability (Brudevold-Newman, 2021; Akresh *et al.,* 2023), and earnings (Duflo, 2001; Khanna, 2020).

In a seminal contribution, Duflo (2001) examines the educational and earnings consequences of a primary school construction programme in Indonesia, the INPRES programme. The paper's main identification strategy is a difference-in-differences framework that leverages cross-cohort variation in exposure to the programme and regional variation in the intensity of the programme, where intensity is measured by the number of schools built across districts. The results suggest that each primary school constructed per 1000 children increased men's education by 0.12 to 0.19 years and wages by 1.5 to 2.7 percent. However, Duflo does not examine whether these gains were also enjoyed by women. Akresh *et al.* (2023) revisit the programme 43 years after its introduction and find it increased women's educational attainment by around 0.23 years but had no measurable impact on their employability.

Unlike the INPRES programme, most schooling expansions in developing countries lack reliable data for impact evaluation (e.g., data on the number of schools built). A growing number of studies circumvent this problem by exploiting information from the pre-programme period. Brudevold-Newman (2021), for instance, examines the effects of a nationwide free secondary education programme in Kenya using a difference-in-differences strategy analogous to that of Duflo (2001), but proxies intensity using a county's primary to secondary school transition rate (the proportion of primary school graduates who enrol in secondary school) in the pre-programme period. Intuitively, the programme could have almost no effect in a county where all primary school graduates were already pursuing secondary education before its introduction, and the largest effects in counties where most students were dropping out after completing primary school. The results indicate that exposure to free secondary education increased female educational attainment by around 0.75 years and induced a shift from agricultural employment towards more skilled jobs. Chicoine (2021), on the other hand, finds no evidence that the free primary education programme in Ethiopia reduced female participation in agriculture. Both studies find the programme in question had no effect on the probability a woman is working.

Research on the effects of the Nigerian UPE programme reveals that it had a positive effect on educational attainment. Leveraging the standard difference-in-differences strategy in the literature, Osili and Long (2008) find the programme increased female schooling by 1.54 years, and Larreguy and Marshall (2017) estimate it increased the probability an individual completed primary school by 18%. Neither of these studies, however, isolate the programme's effects on both men and women, nor do they consider its effect on employability.

II. Background

A. The Nigerian Universal Primary Education Programme

In 1976, faced with rising oil revenues and a desire to curb regional disparities in access to education, Nigeria's military government introduced the UPE programme, which eliminated tuition fees in all public primary schools. To accommodate the expected increase in enrolment, the government introduced plans to construct 150,995 new classrooms at the primary level, and to achieve this objective, the programme mobilized around US$ 900 million to states across the country (Larreguy and Marshall, 2017). Figure 1 illustrates that a significant expansion in public primary school construction, which peaked in 1976, accompanied the programme.

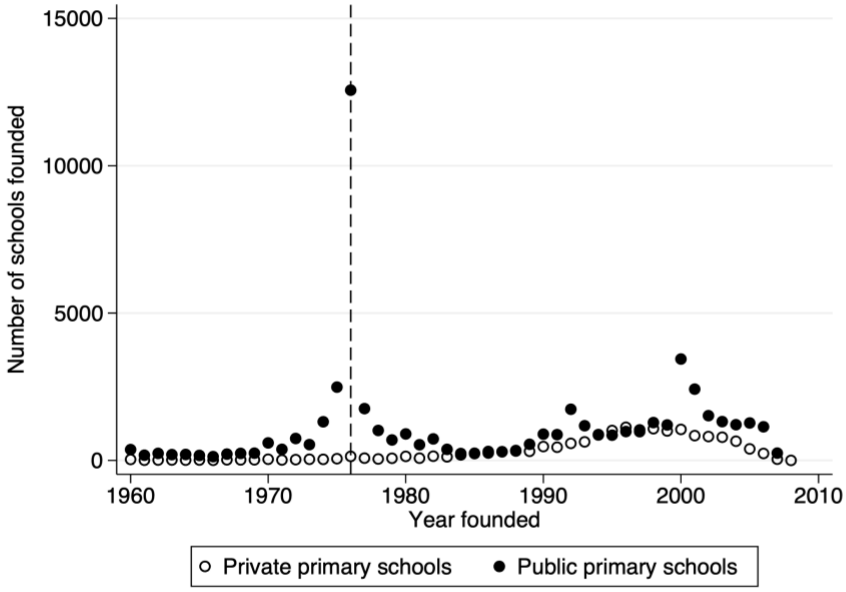

FIGURE 1. NUMBER OF PUBLIC AND PRIVATE SCHOOLS FOUNDED, 1960-2008
Source: Nigeria Primary School Census 2008

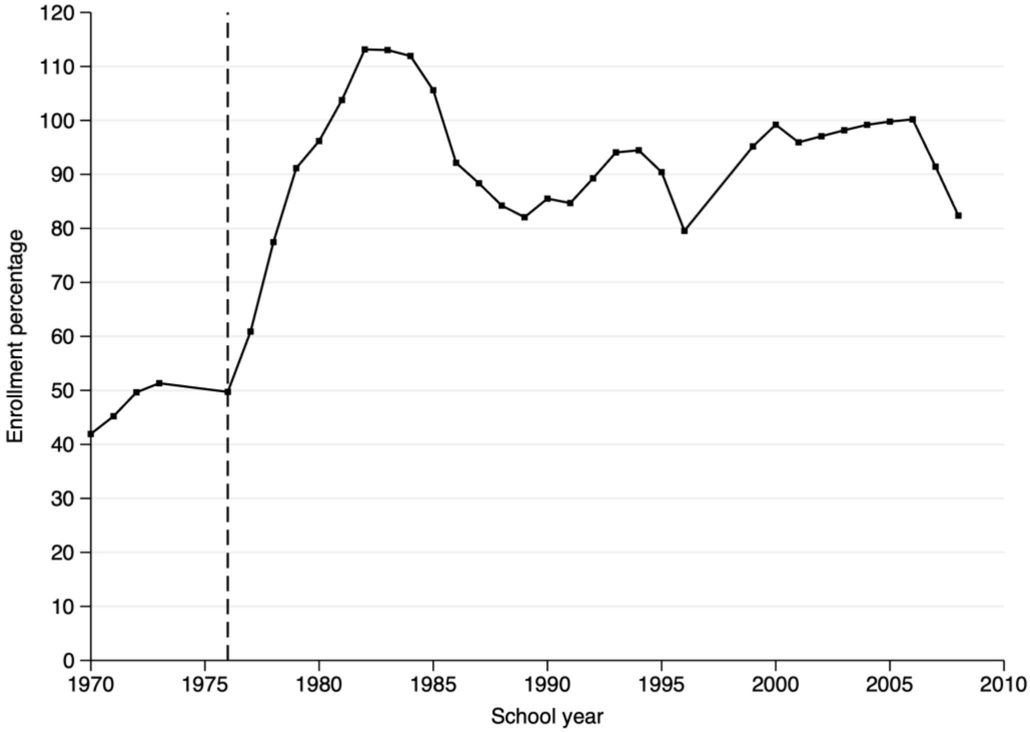

FIGURE 2. PERCENTAGE OF CHILDREN ENROLLED IN PRIMARY SCHOOL
Source: World Development Indicators (2023)

The introduction of the UPE scheme resulted in large increases in enrolment for both boys and girls. Figure 2 illustrates that, between 1976 and 1981, the percentage of children enrolled in primary school more than doubled. Moreover, between 1974 and 1981, the gross primary school enrolment rate for boys rose from 60.3 percent to 136.8 percent, and that for girls increased from 40.3 percent to 104.7 percent (Osili and Long, 2008). These increases in enrolment were not evenly distributed across the country and were the largest in states that had the lowest enrolment rates before the programme (Csapo, 1983).

Due to poor planning, the military government faced some challenges in implementing the programme. In particular, the government underestimated the number of teachers who would be needed to staff new classrooms, and, as a result, schools were often understaffed. In some instances, states remedied this problem by recruiting underqualified teachers with little to no prior training (Csapo, 1983; Asagwara, 1977). The government also underestimated the number of children who would be eligible for the programme, which meant that schools were sometimes unable to accept all the children who wanted to enrol.

A notable aspect of the UPE scheme was that its intensity varied across regions. The programme targeted a 100% gross enrolment rate by 1980 (Larreguy and Marshall, 2017), and thus mobilized resources to states based on the investments required to achieve this target. Because states had different pre-programme enrolment rates, the resources received by each of them varied substantially. These regional differences in programme intensity form the basis of my identification strategy.

Despite its immediate success at increasing enrolment, the UPE programme was only short-lived. The programme was introduced on the premise that the oil boom of the 1970s would continue in the following years, a prediction that was proven wrong with time. Declining oil prices in the late 1970s reduced funding for the programme, ultimately forcing the government to abolish it in 1981. In the years that followed, states responded by reintroducing fees in schools.

III. Empirical Strategy

A. Difference-in-Differences

I identify the effects of the programme on educational and labour market outcomes using a difference-in-differences strategy that exploits cross-cohort and regional variation in exposure to the programme. The first difference of this strategy compares—in states that were more intensely treated by the programme—the schooling and labour market outcomes of cohorts that were exposed to the programme with those of cohorts that did not benefit from the programme. The second difference performs the same comparison in states that were less intensely treated by the programme.

In Nigeria, children typically attend primary school from age 6 to 11. Given that the UPE programme started in 1976 and ended in 1981, its main beneficiaries were children born between 1970 and 1975 (aged 2 - 6 in 1976): these children had the opportunity to start primary school tuition-free. Children who were aged 7 - 11 years in 1976 were also eligible for some free education under the programme, but there is less certainty that they benefited from it. A child aged 9 in 1976, for example, may have already dropped out of school when the programme started, or have never attended school to begin with. I therefore consider individuals born between 1970 and 1975 as being part of the cohorts exposed to the programme. The 1964 cohort was arguably the first cohort unexposed to the programme, since it consists of individuals who were just past primary school age in 1976.[2] I therefore consider individuals born between 1959 and 1964 as having not benefited from the programme.

Following Larreguy and Marshall (2017), I proxy programme intensity using the gap between a state's potential and actual primary school enrolment rate in the pre-programme period. Intuitively, the programme could have little to no effect in a state where all children were already enrolled in primary school before its introduction, and the largest effects in states with the lowest pre-programme primary school enrolment rates. Because men and women have different levels of education, I calculate intensity separately for the two.

For males, I define intensity as one subtract the state's male pre-programme primary school enrolment rate:

$$\text{Male intensity} = (1 - \text{male primary school enrolment rate})$$

For females, I define intensity as one subtract the state's female pre-programme primary school enrolment rate:

$$\text{Female intensity} = (1 - \text{female primary school enrolment rate})$$

Figure 3, which maps for the 36 states and the federal capital territory in Nigeria the estimated male and female programme intensities, illustrates the variation my identification strategy seeks to exploit.

[2] The prevalence of overage enrolment means it is possible that these cohorts were also exposed to the programme. In the cohort analysis (section VI), I explore whether that was the case.

Panel A. Male *Panel B. Female*

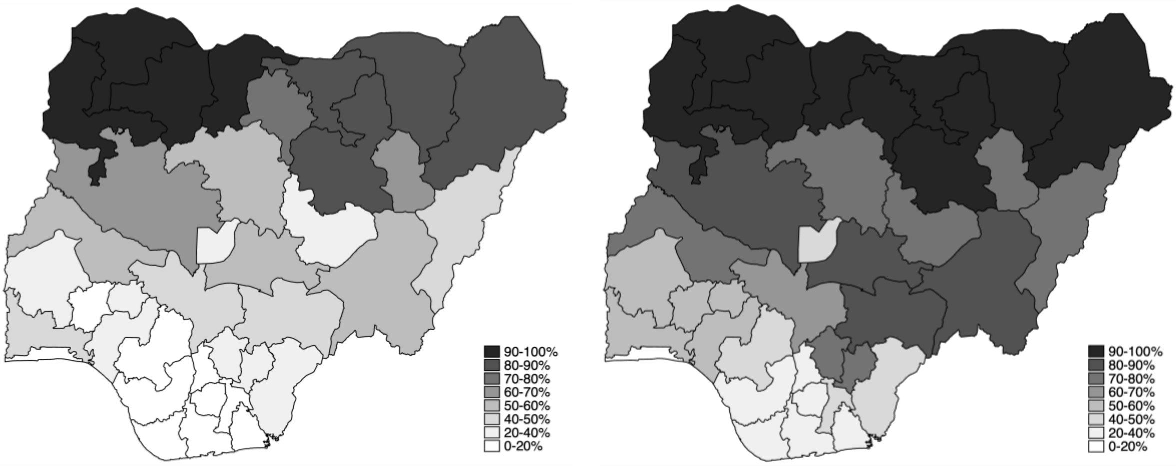

FIGURE 3. SPATIAL VARIATION OF PROGRAMME INTENSITY BY STATE AND GENDER

Source: 2006 Nigeria General Household Survey (GHS)

The basic idea behind the strategy can be illustrated using a two-by-two table. Table 1 presents the means of years of education for men, according to cohort and programme intensity in the state of residence. High-intensity states are those with above-median intensity, and low-intensity states are those with below-median intensity. Column (1) compares the educational attainment of cohorts in high-intensity states that were young enough to benefit from the programme with that of cohorts that were arguably unexposed to the programme. Simply comparing these cohorts, however, does not give the causal effect of the programme, as the gap in educational attainment could reflect unobserved time-varying factors. To isolate the programme's effect, we must compare the cohort difference in high-intensity states with that in low-intensity states (i.e., the difference-in-differences). Column (3) does this comparison and indicates that the cohort difference is larger in high-intensity states (by 0.68 years), which is exactly what we would expect if the programme had any effects.

TABLE 1—MEANS OF EDUCATION BY COHORT AND PROGRAMME INTENSITY

	Programme intensity in state of residence		
	High-intensity	Low-intensity	Difference
	(1)	(2)	(3)
1970-1975 cohorts	5.148 (0.081	9.208 (0.083)	-4.060 (0.117)
1959-1964 cohorts	3.684 0.070)	8.425 (0.087)	-4.741 (0.112)
Difference	1.463 (0.106)	0.782 (0.121)	0.681*** (0.162)

Notes.— The table shows means of years of education for men by cohort and programme intensity in the state of residence. High intensity refers to states with above-median intensity. Low intensity refers to states with below-median intensity. Standard errors are in parentheses. For readability, I display the significance of only the difference-in-differences. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively.

B. Regression Equation

The difference-in-differences specification based on the identification strategy I described in sub-section A is as follows:

$$y_{ijk} = \alpha Intensity_k + \beta\left(UPE_j \times Intensity_k\right) + \rho_k \times Trend_j + \gamma \mathbf{X}_i + \zeta_j + \eta_l + \varepsilon_{ijk} \tag{1}$$

where y_{ijk} is a schooling or labour market outcome of individual i from cohort j in state k; UPE_j is an indicator equal to one if an individual was young enough to benefit from the programme (part of 1970-1975 cohorts) and zero otherwise; $Intensity_k$ is the gender-specific intensity of the programme in state k; and ε_{ijk} is the error term. I control for time-invariant differences between cohorts and states by including cohort and state fixed effects, denoted by ζ_j and η_k, respectively. Following the conventions in the difference-in-differences literature, I control for state-specific linear trends, $\rho_k \times Trend_j$, which account for potential differential pre-programme trends among states. I increase the efficiency of my estimates by including a set of religion dummies and a rural-urban dummy in $\mathbf{X}_i$. Following Larreguy and Marshall (2017), I account for spatial clustering in the programme's intensity by clustering standard errors at the state level throughout. The parameter of interest is β, which is the difference-in-differences estimate of the effect of the programme on schooling and labour market outcomes. I estimate specification (1) separately for men and women.

C. Key Identifying Assumptions

1. Parallel Trends in Schooling and Labour Market Outcomes

The validity of the identification strategy relies on the assumption that, in the absence of the UPE programme, schooling and labour market outcomes in high-intensity and low-intensity states would have followed the same trend. This assumption is violated if, for example, even without the introduction of the programme, individuals in high-intensity states would have experienced larger increases in educational attainment than those in low-intensity states. In such a case, the estimated programme effect would be biased.

Because we do not observe schooling and labour market outcomes in the absence of the programme (i.e., the counterfactual), we cannot conclusively validate this assumption. Nonetheless, it is useful to visually inspect the pre-programme trends in the outcomes of interest for the two groups of states. Intuitively, if schooling and labour market outcomes before the introduction of the programme in low- and high-intensity states moved parallel to each other, then we can plausibly assume that they would have continued to do so had the programme not been introduced. Thus, we can use schooling and labour market outcomes in low-intensity states as a counterfactual for those in high-intensity states.

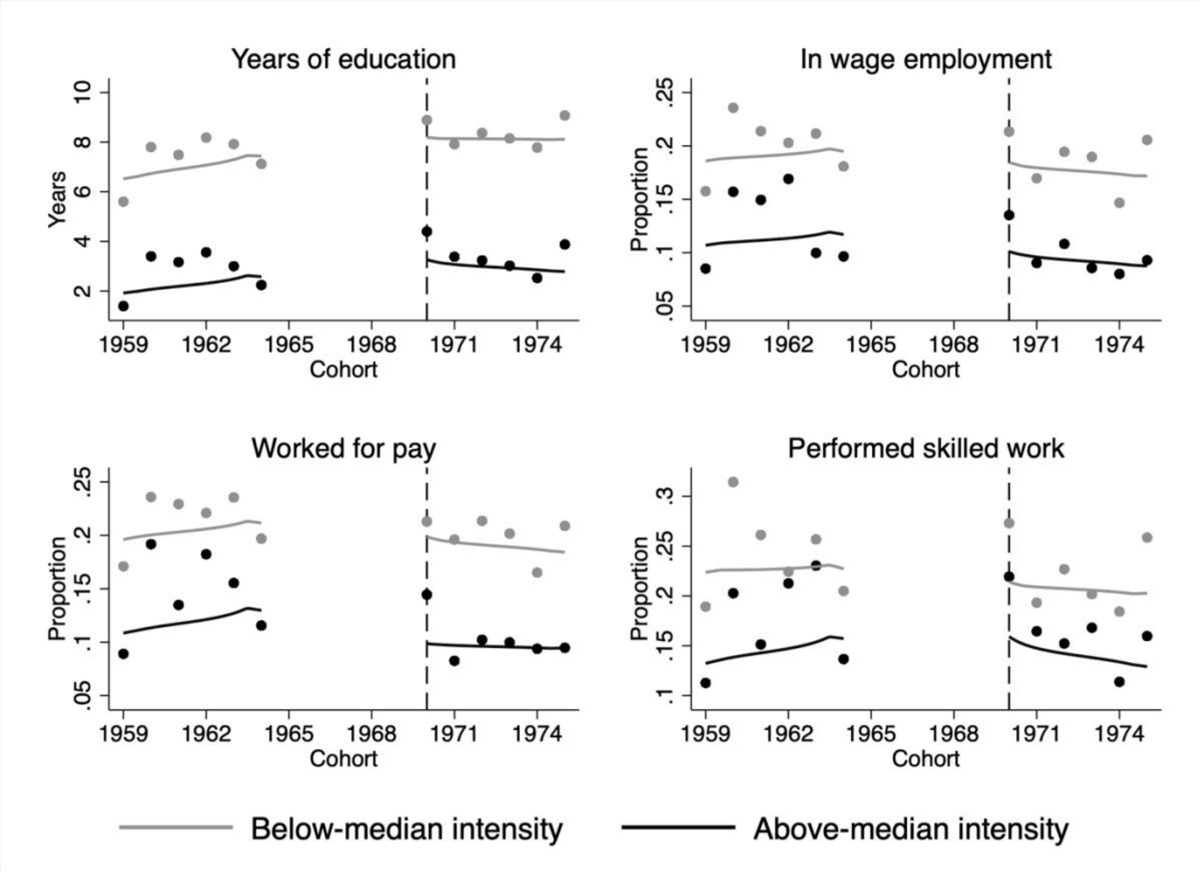

FIGURE 4. TRENDS IN SCHOOLING AND LABOUR MARKET OUTCOMES

NOTE.—In the first panel, each point shows the average years of education in states with above-median and below-median intensities. I overlay local polynomials to show the trends in educational attainment across treatment intensity. Above-median intensity and below-median intensity denote respondents living in states with above-median intensity and below-median intensity, respectively. The vertical dashed line is on the 1970 cohort. The sample includes all individuals who are part of the cohorts that define exposure to the programme (1959-1964 and 1970-1975 cohorts).

In Figure 4, I plot the trends in years of schooling and three labour market outcomes for states with below-median intensity and those with above-median intensity. The figure shows similar pre-programme trends in these outcomes, lending support to my identification strategy. Moreover, since I include state-specific linear trends in all specifications, any unobserved differential trends among states should be accounted for in estimation. Nevertheless, in robustness checks, I formally test the plausibility of this assumption.

2. No Alternative Policy Changes or Shocks

The identification strategy will provide unbiased estimates of the treatment effect only if no omitted policy changes or shocks correlate with the programme, as the estimator might confound these changes with the effect of the programme. During the period of study, the Nigerian labour market underwent some changes worth highlighting. In particular, employment in the public sector and the wages of public servants increased (Osili and Long, 2008). These labour market reforms could have encouraged girls and boys to become educated or to enter the labour market, which would compromise identification. Following Duflo (2001), I account for these effects by presenting specifications that control for interactions between cohort dummies and the 1976 state unemployment rate.

V. Data

A. Dataset

I use data from the 2009 Harmonised Nigeria Living Standard Survey (HNLSS), a nationally representative cross-section of 77,390 households in Nigeria.[3] The HNLSS collects rich information on educational attainment and employment outcomes, making it particularly suited for this study. The survey has the added advantage that, because it was conducted 33 years after the programme's introduction, it allows me to study the programme's beneficiaries in their prime working age.

---

[3] I am grateful to Professor Horacio Larreguy from the Instituto Tecnológico Autónomo de México (ITAM) for facilitating access to the HNLSS data. The data are downloadable at: https://www.nigerianstat.gov.ng/nada/index.php/catalog/38

A limitation of the HNLSS, however, is that it reports the earnings of only individuals working in wage jobs, who make up less than 15 percent of the sample. Given the unreliability of the earnings data, my labour market analysis does not consider the productivity effects of the programme. Instead, I focus on its effects on employability.

Because I am primarily interested in how the UPE programme affected educational and employment outcomes, I restrict the sample to the cohorts that define exposure to the programme: the 1970-1975 and 1959-1964 cohorts. To isolate the programme's effect on men and women, I divide this sample into a male sample and a female sample (see Table 1 for descriptive statistics).

B. Outcome Variables

Schooling outcomes.—I use eight schooling outcomes: *Years of Education, Completed primary, Completed junior secondary, Completed senior secondary, Enrolled in tertiary, Can read in English, Can write in English, and Can perform written arithmetic.* All these are dummy variables except *Years of education. Enrolled in tertiary* equals one if an individual has ever enrolled in a post-secondary institution. *Can read in English* and *Can write in English* equal one if an individual can read or write a letter in English. *Can perform written arithmetic* equals one if an individual can perform simple written calculations. The definitions of the other variables are clear from their names.

Labour market outcomes.—I use five labour market outcomes: *Worked for pay*, *In wage employment, Performed skilled work, Worked in Public sector,* and *Worked in agriculture*. All these are dummy variables. *Worked for pay* equals one if an individual worked for pay in the past seven days. She was in wage employment if she had a job that sets explicit or implicit employment contracts and pays a salary independent of output. She performed skilled work if she held a managerial, professional, or technical position in the past seven days. She worked in the public sector if she spent most time working as a government employee in the past twelve months. She participated in agriculture if she spent most time in self-employed agricultural work in the past twelve months.

Table 1 presents descriptive statistics for the male and female samples. Individuals in the male sample are 42 years old and have 6 years of schooling on average. About 55 percent, 36 percent, and 33 percent completed primary, junior secondary, and senior secondary schools, respectively. About one in two can read in English, write in English, and perform simple written arithmetic. About one in five were in wage employment, worked for pay, or performed skilled work. Most worked in agriculture, and about one in ten worked in the public sector.

Individuals in the female sample are 41 years old and have 4 years of schooling on average. About 40 percent, 21 percent, and 18 percent completed primary, junior secondary, and senior secondary schools, respectively. About one in three can read in English, write in English, or perform simple written arithmetic. Around one in ten were in wage employment or worked for pay. One in five performed skill work, less than one in ten worked in the public sector, and one in three worked in agriculture.

C. Computation and Measurement of Intensity

As highlighted in section IV, I proxy programme intensity using the gap between a state's actual and potential primary school enrolment rate in the pre-programme period. I calculate the enrolment rates using data from the 2006 General Household Survey (GHS).[4] The GHS is similar to the HNLSS in that it is nationally representative and collects information on educational attainment and socioeconomic status.

Following Larreguy and Marshall (2017), I calculate for each state the proportion of all males or females born between 1955 and 1964 who had completed primary school. Intuitively, individuals born between 1955 and 1964 make up the first ten cohorts arguably unexposed to the programme, and thus can be used to estimate the pre-programme primary school enrolment rates.[5] More precisely, I calculate the enrolment rate for gender l in state k as follows:

$$Primary\ school\ enrolment\ rate_{kl} = \frac{\sum_{j=1955}^{1964} Completed_{jkl}}{\sum_{j=1955}^{1964} Total_{jkl}} \tag{2}$$

where $Completed_{jkl}$ is the number of individuals of gender l, from cohort j, in state k completing primary school, and $Total_{jkl}$ is the number of all individuals of gender l, from cohort j, in state k.

I match these state-level data to individual-level data from the HNLSS. The HNLSS reports an individual's state of residence but not the state in which she attended primary school. Thus, in the matching process, I assume respondents were educated in the state in which they currently reside. A threat to this assumption is selective migration. For example, if educated youths migrated to states that were more intensely treated by the programme in search of job opportunities, the estimator could confound these movements with a programme effect.

[4] The data are downloadable at: https://www.nigerianstat.gov.ng/nada/index.php/catalog/22

[5] The choice of ten cohorts is to ensure I am not calculating enrolment rates from a small number of observations.

TABLE 2—DESCRIPTIVE STATISTICS

	Male	Female
Panel A. Schooling outcomes		
Years of education	6.13	4.07
	(6.25)	(5.36)
Completed primary	0.55	0.40
	(0.50)	(0.49)
Completed junior secondary	0.36	0.21
	(0.48)	(0.41)
Completed senior secondary	0.33	0.18
	(0.47)	(0.39)
Can read in English	0.51	0.36
	(0.50)	(0.48)
Can write in English	0.49	0.34
	(0.50)	(0.47)
Can perform written arithmetic	0.51	0.37
	(0.50)	(0.48)
Panel B. Labour market outcomes		
Worked for pay	0.21	0.11
	(0.41)	(0.31)
In wage employment	0.18	0.10
	(0.39)	(0.29)
Performed skilled work	0.19	0.19
	(0.39)	(0.39)
Worked in the public sector	0.11	0.06
	(0.32)	(0.24)
Worked in agriculture	0.60	0.35
	(0.49)	(0.48)
Panel C. Control variables		
Lives in rural area	0.75	0.74
	(0.44)	(0.44)
Age	42.38	41.56
	(6.02)	(6.08)
Christian	0.48	0.52
	(0.50)	(0.50)
Muslim	0.51	0.46
	(0.50)	0.50)
Other religion	0.02	0.02
	(0.13)	(0.13)

Notes.— The table presents the means and standard deviations for males and females. The number in each cell is the mean; the figures in parentheses are standard deviations. The male and female samples include individuals in the cohorts that define exposure to the UPE programme (1959-1964 and 1970-1975 cohorts). All the variables are dummy variables except Years of education and Age. The number of observations of most variables varies from 20400 to 21852 except Worked for pay and Performed skilled work. For these two variables, the number of observations varies from 11332 to 12759.

In the context of Nigeria, however, migration is unlikely to threaten identification for two reasons. First, migration is uncommon and happens mostly within states (Osili and Long, 2008; Oyelere, 2010; Larreguy and Marshall, 2017). In the HNLSS sample, 61% of respondents report always living at their current location; among those who report moving, 40% report migrating within the same state. Second, existing studies find no evidence that migrants and non-migrants differ in their educational outcomes (Osili and Long, 2008). Nevertheless, in robustness checks, I examine how addressing potential selective migration affects my results.

IV. Results

In this section, I present estimates of the programme's effects on educational attainment, learning outcomes, and labour market outcomes.

A. Effects on Schooling Outcomes

Table 3 presents the estimated impacts on educational attainment (Panel A for men, and Panel B for women). Each estimate represents the effect of moving from the lowest to the highest intensity state during the programme period. To get estimated effects comparable to those in existing studies, it is useful to convert the estimates reported in tables into effects at mean intensity.[6]

Almost all the estimates in Panel A are positive and statistically significant, which suggests that the programme induced boys in states with low pre-programme primary school enrolment rates to attend school. The estimate in column (1) indicates that, at the mean male intensity value of 0.40, the programme increased men's educational attainment by around one year. Controlling for the effects of employment initiatives that could correlate with the UPE programme, I also find it increased male educational attainment by around one year (column (2)). Economically,

[6] To get effects at mean intensity, I multiply the reported estimates by the mean intensity in the samples. In the male sample, the mean intensity is 0.40; in the female sample, it is 0.70.

these effects are large. Given that the average male educational attainment is around 6 years, the estimate in column (2) amounts to a 15.8 percent increase in educational attainment.

TABLE 3—EFFECTS OF UPE ON EDUCATIONAL ATTAINMENT

	Years of education		Completed primary		Completed junior secondary		Completed senior secondary		Enrolled in tertiary	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A. Male										
UPE x Intensity	2.619***	2.421***	0.251***	0.189**	0.204**	0.193**	0.185**	0.182**	0.039	0.061
	(0.674)	(0.780)	(0.073)	(0.072)	(0.079)	(0.086)	(0.079)	(0.089)	(0.044)	(0.051)
Year of birth × 1976 unemployment rate		✓		✓		✓		✓		✓
Observations	21426	21426	21493	21493	21493	21493	21493	21493	21493	21493
Outcome mean	6.135	6.135	0.553	0.553	0.356	0.356	0.327	0.327	0.143	0.143
Adjusted R^2	0.276	0.276	0.295	0.295	0.205	0.205	0.194	0.193	0.086	0.086
Panel B. Female										
UPE x Intensity	1.983*	1.561	0.202***	0.211**	0.044	-0.006	0.004	-0.046	0.065	0.009
	(1.145)	(1.393)	(0.069)	(0.085)	(0.103)	(0.124)	(0.102)	(0.124)	(0.056)	(0.053)
Year of birth × 1976 unemployment rate		✓		✓		✓		✓		✓
Observations	21734	21734	21750	21750	21750	21750	21750	21750	21750	21750
Outcome mean	4.070	4.070	0.401	0.401	0.208	0.208	0.182	0.182	0.075	0.075
Adjusted R^2	0.330	0.331	0.337	0.337	0.214	0.214	0.189	0.190	0.078	0.078

NOTE.—The table shows the results of specification (1) estimated for males and females. Each column is a different regression. The dependent variables are listed in column headings. All the variables are dummy variables except for Years of education. All specifications control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, and a rural-urban dummy. Regressions in even-numbered columns also control for interactions between cohort dummies and the 1976 unemployment rate in an individual's state of residence. Standard errors, reported in parentheses, are clustered at the state level. The number of observations varies across specifications due to data availability. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively.

Although the UPE programme abolished fees only in primary schools, the results in columns (3) and (4) suggest that it increased male educational attainment at both the primary and secondary levels: at mean intensity, the programme increased the probability a man completed primary school by 7.6 percentage points (13.7 percent), junior secondary school by 7.7 percentage points (21.7 percent), and senior secondary school by 7.3 percentage points (22.3 percent). These estimates are robust to accounting for employment initiatives that could correlate with the UPE programme. However, I find no evidence that the programme increased male participation in tertiary education (columns (9) and (10)).

Panel B indicates that, at the mean female intensity value of 0.70, the programme increased female educational attainment by around one year (Column (2)). This effect is economically large but imprecisely estimated. The impact on primary school completion (column (4)) is estimated with greater precision and suggests that, at mean intensity, the programme increased the probability a woman completed primary school by around 14.8 percentage points (36.8 percent). However, I find no evidence that the programme increased female educational attainment beyond the primary level.[7] If anything, estimates of the effects on junior secondary school and senior secondary school completion are negative, and those on enrolment in tertiary education are small.

Table 4 presents estimates of the programme's effects on literacy and arithmetic skills. Column (5) of Panel A suggests that the programme increased the probability a man can perform written arithmetic, but controlling for the effects of employment initiatives that could correlate with the UPE scheme renders this estimate statistically insignificant. Column (2) of Panel B suggests that, at mean intensity, the programme increased the probability a woman can read in English by 10.1 percentage points (28.6 percent), but this estimate is significant only at the ten percent level.

[7] I also estimate the programme's impact on female enrolment—rather than on completion—in secondary school (results not presented in paper) and find no significant effect.

TABLE 4—EFFECTS OF UPE ON LITERACY AND ARITHMETIC SKILLS

	Can read in English		Can write in English		Can perform written arithmetic	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Male						
UPE x Intensity	0.125*	0.083	0.128*	0.088	0.169**	0.127
	(0.067)	(0.079)	(0.067)	(0.084)	(0.065)	(0.083)
Year of birth × 1976 unemployment rate		✓		✓		✓
Observations	21465	21465	21465	21465	21465	21465
Outcome mean	0.514	0.514	0.487	0.487	0.515	0.515
Adjusted R^2	0.320	0.320	0.307	0.307	0.296	0.296
Panel B. Female						
UPE x Intensity	0.169**	0.148*	0.185**	0.132	0.121	0.097
	(0.072)	(0.083)	(0.082)	(0.091)	(0.076)	(0.080)
Year of birth × 1976 unemployment rate		✓		✓		✓
Observations	21669	21669	21669	21669	21669	21669
Outcome mean	0.362	0.362	0.336	0.336	0.371	0.371
Adjusted R^2	0.330	0.330	0.308	0.308	0.311	0.311

Notes.—The table shows the results of specification (1) estimated for males and females. Each column is a different regression. All regressions are estimated using OLS estimators. The dependent variables are listed in column headings. All the variables are dummy variables. All regressions control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, and a rural-urban dummy. Specifications in even-numbered columns also control for interactions between cohort dummies and the 1976 unemployment rate in the state of residence. Standard errors, reported in parentheses, are clustered at the state level. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively.

B. *Effects on Labour Market Outcomes*

Table 5 presents estimates of the programme's effects on employment outcomes (Panel A for men, and Panel B for women). I find no evidence that the programme increased male participation in wage employment or paid work. The programme also does not seem to have increased male participation skilled work and the public sector (columns (5)-(8)). However, I find it reduced male participation in agriculture by around 6.7 percentage points (11.1 percent) at mean intensity.

TABLE 5—EFFECTS OF UPE ON LABOUR MARKET OUTCOMES

	In wage employment		Worked for pay		Performed skilled work		Worked in public sector		Worked in agriculture	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel B. Male										
UPE x Intensity	-0.037	0.038	-0.179**	-0.132	0.029	0.140	-0.060	0.010	-0.163**	-0.168**
	(0.082)	(0.088)	((0.082)	(0.092)	(0.076)	(0.092)	(0.043)	(0.044)	(0.061)	(0.062)
Year of birth × 1976 unemployment rate		✓		✓		✓		✓		✓
Observations	20400	20400	13844	13844	12755	12755	20400	20400	20400	20400
Outcome mean	0.177	0.177	0.207	0.207	0.189	0.189	0.113	0.113	0.605	0.605
Adjusted R^2	0.077	0.078	0.079	0.080	0.069	0.071	0.066	0.067	0.291	0.292
Panel A. Female										
UPE x Intensity	0.073	0.054	0.050	0.015	-0.041	-0.046	0.055	0.042	0.059	0.017
	(0.053)	(0.055)	(0.071)	(0.070)	(0.100)	(0.104)	(0.042)	(0.042)	(0.097)	(0.107)
Year of birth × 1976 unemployment rate		✓		✓		✓		✓		✓
Observations	18664	18664	12759	12759	11332	11332	18664	18664	18664	18664
Outcome mean	0.086	0.086	0.110	0.110	0.187	0.187	0.061	0.061	0.354	0.354
Adjusted R^2	0.069	0.069	0.063	0.063	0.088	0.089	0.069	0.070	0.189	0.192

Notes.—The table shows the results of specification (1) estimated for males and females. Each column is a different specification. The dependent variables are in column headings. All the variables are dummy variables. All specifications control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, and a rural-urban dummy. Specifications in even-numbered columns also control for interactions between cohort dummies and the 1976 unemployment rate in the state of residence. Standard errors, reported in parentheses, are clustered at the state level. The number of observations varies across specifications due to data availability. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively.

I find no evidence that the programme improved female employability (Panel B): the estimates of the effect on female participation in wage employment (columns (1) and (2)) and the public sector (columns (7) and (8)) are statistically insignificant at the conventional level of significance.

C. *Effect by Cohort*

The results presented in the previous sub-sections indicate that the programme increased male educational attainment, reduced male participation in agriculture, and increased the probability a woman completed primary school. Even after accounting for differential pre-programme trends among states and controlling for the effects of employment initiatives that could correlate with the UPE scheme, I cannot rule out the possibility that the estimator is capturing unobserved time-varying factors or mean reversions in schooling and employment outcomes that would have happened even in the absence of the programme.

Following Duflo (2001), I explore the extent to which my results are spurious by conducting a cohort analysis, which has the advantage of isolating the programme's effect on each birth cohort. To conduct this analysis, I expand the sample to all individuals born between 1959 and 1975. I then estimate the following specification:

$$y_{ijk} = \alpha_1 Intensity_k + \sum_{j=1960}^{1975} \left(Born_{ij} \times Intensity_k\right)\beta_{1j} + \sum_{j=1960}^{1975} \left(Born_{ij} \times Unem_k\right)\lambda_{1j} + \rho_l \times Trend_j + \gamma \mathbf{X}_i + \zeta_j + \eta_l + \varepsilon_{ijkl} \quad (2)$$

where y_{ijk} is a schooling or labour market outcome of individual i from cohort j in state k; $Intensity_k$ is the gender-specific intensity of the programme in state k (as in specification(1)); ζ_j and η_l are cohort and state fixed effects, respectively; $\rho_l \times Trend_j$ are state-specific linear trends; $\mathbf{X}_i$ is a vector of control variables, which includes a set of religion dummies and a rural-urban dummy; and ε_{ijk} is the error term. $Born_{ij}$ is a dummy equal to one if individual i is part of cohort *j*. The reference group for this dummy is the 1959 cohort. Following the result from the previous sub-sections that the basic specification might be biased, I also control for interactions between cohort dummies and the 1976 state unemployment rate, $\sum_{j=1960}^{1975}\left(Born_{ij} \times Unem_k\right)$, where $Unem_k$ is the 1976 unemployment rate in state k. The coefficients β_{1j} measure the programme's impact on each birth cohort between the years 1960 and 1975.

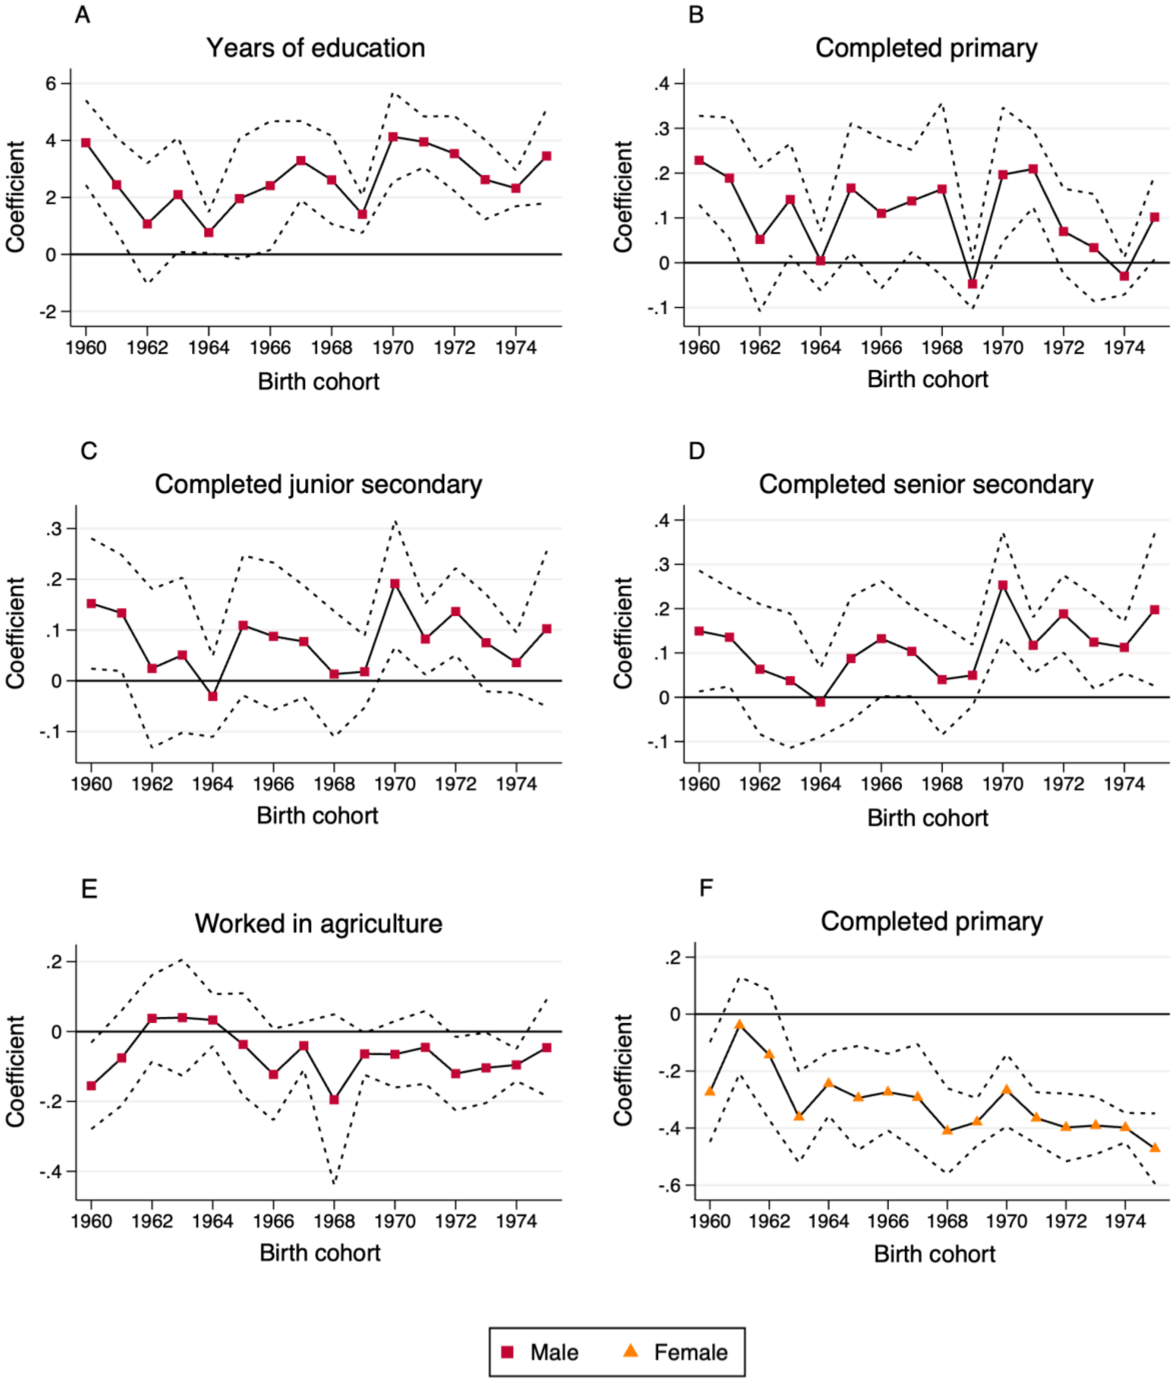

FIGURE 4. COEFFICIENTS OF INTERACTIONS BETWEEN YEAR OF BIRTH AND PROGRAMME INTENSITY IN THE STATE OF RESIDENCE WITH NINETY-FIFTH PERCENT CONFIDENCE INTERVAL BANDS

Notes.—The figure shows estimates of the effects of the programme on each birth cohort. Each panel represents a different regression. The title of each panel denotes the dependent variable. The omitted cohort (control cohort) in each regression is the 1959 cohort.

Figure 4 plots estimates of these coefficients. Motivated by the results from sub-sections A and

B, I focus on the effects on educational attainment and male participation in agriculture. The horizontal axis of each panel denotes the birth cohort, and the vertical axis measures the effect of the programme on each cohort (the coefficients on the interactions, β_{1j}). The solid lines connect my point estimates, and the dashed lines represent 95 percent confidence intervals. Red squares and yellow triangles represent point estimates for men and women, respectively.

Several features of these graphs stand out. First, in Panel A, the estimated effect on male years of schooling jumps up, starting with the 1970 cohort—the cohort that could attend primary school for the longest duration under the programme. The effect also declines with younger cohorts, reflecting the fact that they benefited less from the programme. The estimated effects on primary school completion in Panel B exhibit a similar pattern. This pattern is clearer in Panels C and D, which illustrate the estimated effects on male junior secondary school and senior secondary school completion.

Panels A and B also reveal that the effects on years of education and primary school completion are also positive and statistically significant for some of the 1965-1969 cohorts (those that I omit from my difference-in-differences analysis), which is not surprising, given these individuals were of primary school age when the programme was implemented. A final implication of these results is that the 1959-1964 cohorts may not be perfect control cohorts, as there is evidence of exposure among men born in 1960 and 1961, which is most likely due to overage enrolment. Panel E illustrates the effects on male participation in agricultural work. The negative effects on participation in agriculture are concentrated among men born after 1969, who were the main beneficiaries of the programme. Altogether, the cohort analysis suggests that the difference-in-differences estimates of the programme's effects on male outcomes are not spurious.

The results for women in Panel F, however, seem to contradict those presented in the previous sub-sections: the estimated effects on primary school completion are negative for almost all cohorts and decrease as we move towards younger cohorts, which suggests that, during the period of study, the proportion of women completing primary school in low-intensity states increased by more than that in high-intensity states. Many factors could explain these results. Drawing on the conceptual framework introduced in section II, I next explore the role of cultural factors.

D. The Role of Cultural Factors

In this section, I examine the extent to which cultural factors may drive the negative results for women. At the time when the programme was launched, cultural prohibitions against educating girls were prevalent in Nigeria, especially among the Muslim population (Clarke, 1978; Csapo, 1981; Csapo, 1983). The conceptual framework introduced in section II predicts that such forces should counter the expected positive impact of the UPE scheme on female education and employment.

Motivated by evidence that Muslim parents opposed `Western' education as a means of educating their girls (Clarke, 1978), I divide the female sample into two subsamples: one Muslim and one non-Muslim. The former consists of female respondents reporting Islam as their religion, and the latter consists of women subscribing to other religions. I re-estimate specification (1) for these two subsamples separately. As in the previous sub-section, I also control for interactions between cohort dummies and the 1976 unemployment rate in all specifications.

TABLE 6—EFFECTS OF UPE ON EDUCATIONAL ATTAINMENT

	Non-Muslim					Muslim				
	Years of Education	Completed Primary	Completed junior secondary	Completed senior secondary	Enrolled in tertiary	Years of Education	Completed Primary	Completed junior secondary	Completed senior secondary	Enrolled in tertiary
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
UPE x Intensity	3.274*	0.348**	0.114	0.049	0.039	-1.052	0.073	-0.237	-0.256	0.044
	(1.845)	(0.145)	(0.137)	(0.134)	(0.089)	(2.397)	(0.138)	(0.271)	(0.268)	(0.065)
Observations	11660	11673	11673	11673	11673	10074	10077	10077	10077	10077
Outcome mean	6.088	0.591	0.315	0.274	0.117	1.736	0.182	0.085	0.075	0.026
Adjusted R^2	0.202	0.183	0.159	0.143	0.064	0.218	0.248	0.128	0.119	0.037

Notes.—The table shows the results of specification (1) estimated for Muslim and Non-Muslim women. Each column is a different specification. The dependent variables are in column headings. All the variables are dummy variables except Years of education. All specifications control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, a rural-urban dummy, and interactions between cohort dummies and the 1976 unemployment rate in the state of residence. Standard errors, reported in parentheses, are clustered at the state level. The number of observations varies across specifications due to data availability. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively. The sample includes individuals part of the cohorts that define exposure to the UPE programme (the 1959-1964 and 1970-1975 cohorts).

TABLE 7—EFFECTS OF UPE ON LITERACY AND ARITHMETIC SKILLS

	Non-Muslim			Muslim		
	Can read in English	Can write in English	Can perform written arithmetic	Can read in English	Can write in English	Can perform written arithmetic
	(1)	(2)	(3)	(4)	(5)	(6)
UPE x Intensity	0.342***	0.326***	0.227**	-0.103	-0.222	-0.055
	(0.119)	(0.117)	(0.111)	(0.142)	(0.161)	(0.175)
Observations	11617	11617	11617	10052	10052	10052
Outcome mean	0.551	0.517	0.550	0.143	0.127	0.165
Adjusted R^2	0.185	0.172	0.178	0.202	0.174	0.206

Notes.—The table shows the results of specification (1) estimated for Muslim and Non-Muslim women. Each column is a different specification. The dependent variables are in column headings. All the variables are dummy variables except Years of education. All specifications control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, a rural-urban dummy, and interactions between cohort dummies and the 1976 unemployment rate in the state of residence. Standard errors, reported in parentheses, are clustered at the state level. The number of observations varies across specifications due to data availability. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively. The sample includes individuals part of the cohorts that define exposure to the UPE programme (the 1959-1964 and 1970-1975 cohorts).

TABLE 8—EFFECTS OF UPE ON EMPLOYMENT STATUS AND SECTOR OF EMPLOYMENT

	Non-Muslim					Muslim				
	In wage employment	Worked for pay	Performed skilled worked	Worked in Public sector	Worked in agriculture	In wage employment	Worked for pay	Performed skilled work	Worked in Public sector	Worked in agriculture
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
UPE x Intensity	0.009	-0.009	-0.065	0.044	-0.224	0.139	0.050	0.152	0.002	0.070
	(0.071)	(0.124)	(0.186)	(0.058)	(0.135)	(0.109)	(0.143)	(0.174)	(0.083)	(0.152)
Observations	10911	7865	7122	10911	10911	7753	4894	4210	7753	7753
Outcome mean	0.123	0.143	0.215	0.089	0.452	0.042	0.057	0.138	0.021	0.216
Adjusted R^2	0.065	0.070	0.099	0.067	0.067	0.191	0.076	0.044	0.108	0.039

Notes.—The table shows the results of specification (1) estimated for Muslim and Non-Muslim women. Each column is a different specification. The dependent variables are in column headings. All the variables are dummy variables except Years of education. All specifications control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, a rural-urban dummy, and interactions between cohort dummies and the 1976 unemployment rate in the state of residence. Standard errors, reported in parentheses, are clustered at the state level. The number of observations varies across specifications due to data availability. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively. The sample includes individuals part of the cohorts that define exposure to the UPE programme (the 1959-1964 and 1970-1975 cohorts).

Table 6 presents estimates of the programme's effects on educational attainment. The positive effects on female education presented in section C seem to be concentrated among non-Muslims (Columns (1) and (2)): at the mean non-Muslim intensity value of 0.55, the programme increased the educational attainment of non-Muslim women by 1.8 years (29.6 percent). However, this effect is imprecisely estimated and statistically significant only at the ten percent level. I also find the programme increased the probability a non-Muslim woman completed primary school by 19.1 percentage points (32.3 percent). The effects on the educational attainment of Muslim women are statistically insignificant and mostly negative.

Table 7 presents estimates of the impacts on learning outcomes. Unlike in the previous analyses, I find the programme improved female literacy and arithmetic skills, but these effects seem to be concentrated among non-Muslims. At mean intensity, the programme increased the probability a non-Muslim woman can read in English, write in English, and Perform written calculations by 18.8, 17.9, and 12.5 percentage points, respectively (columns (1) - (3)). The effects on non-Muslims are always negative and statistically insignificant.

Table 8 presents estimates of the programme's effects on employment outcomes. These results are consistent with the finding from previous analyses that the programme had no statistically significant effect on female employment. The result in column (5) indicates that, at mean intensity, the programme reduced female participation in agricultural work by 12.3 percentage points (42.2 percent). Though economically large, this estimate is statistically insignificant.

To ensure that these results are not spurious, I replicate the cohort analysis for the two subsamples. Figure 5 presents the results. Motivated by the difference-in-differences results, I focus on the effects on primary school completion and learning outcomes. The left and right panels illustrate the results for non-Muslim women and Muslim women, respectively. The panels for non-Muslim women indicate that the positive effects are concentrated among the 1970-1975 cohorts, suggesting the main difference-in-differences estimates are not spurious. However, in Panel D, although the treatment effect jumps up, starting with the 1970 cohort, the estimates are not statistically significant at conventional levels.

In the Muslim subgroup (right panels), the estimated effects are negative for most cohorts and decline with younger cohorts. It also appears these effects were declining in the pre-programme period (for the 1960-1964 cohorts), suggesting that we cannot attribute these declining effects to the UPE programme. A thorough analysis of why schooling outcomes of Muslim women in high-intensity states were declining relative to those in low-intensity states would be beyond the scope of this paper. However, the fact that the estimated programme effects for the 1970-1975 cohorts in the Muslim sample failed to increase as in the non-Muslim sample may provide suggestive evidence that prohibition towards educating girls may have undermined the accumulation of human capital among Muslim girls.

V. Robustness Checks

In this section, I test the robustness of the paper's main findings by addressing selective migration, testing the plausibility of the parallel trends assumption, and conducting placebo tests. Following evidence of heterogeneity in the effects of the programme on Muslim and non-Muslim women, I conduct robustness checks separately for these two subgroups.

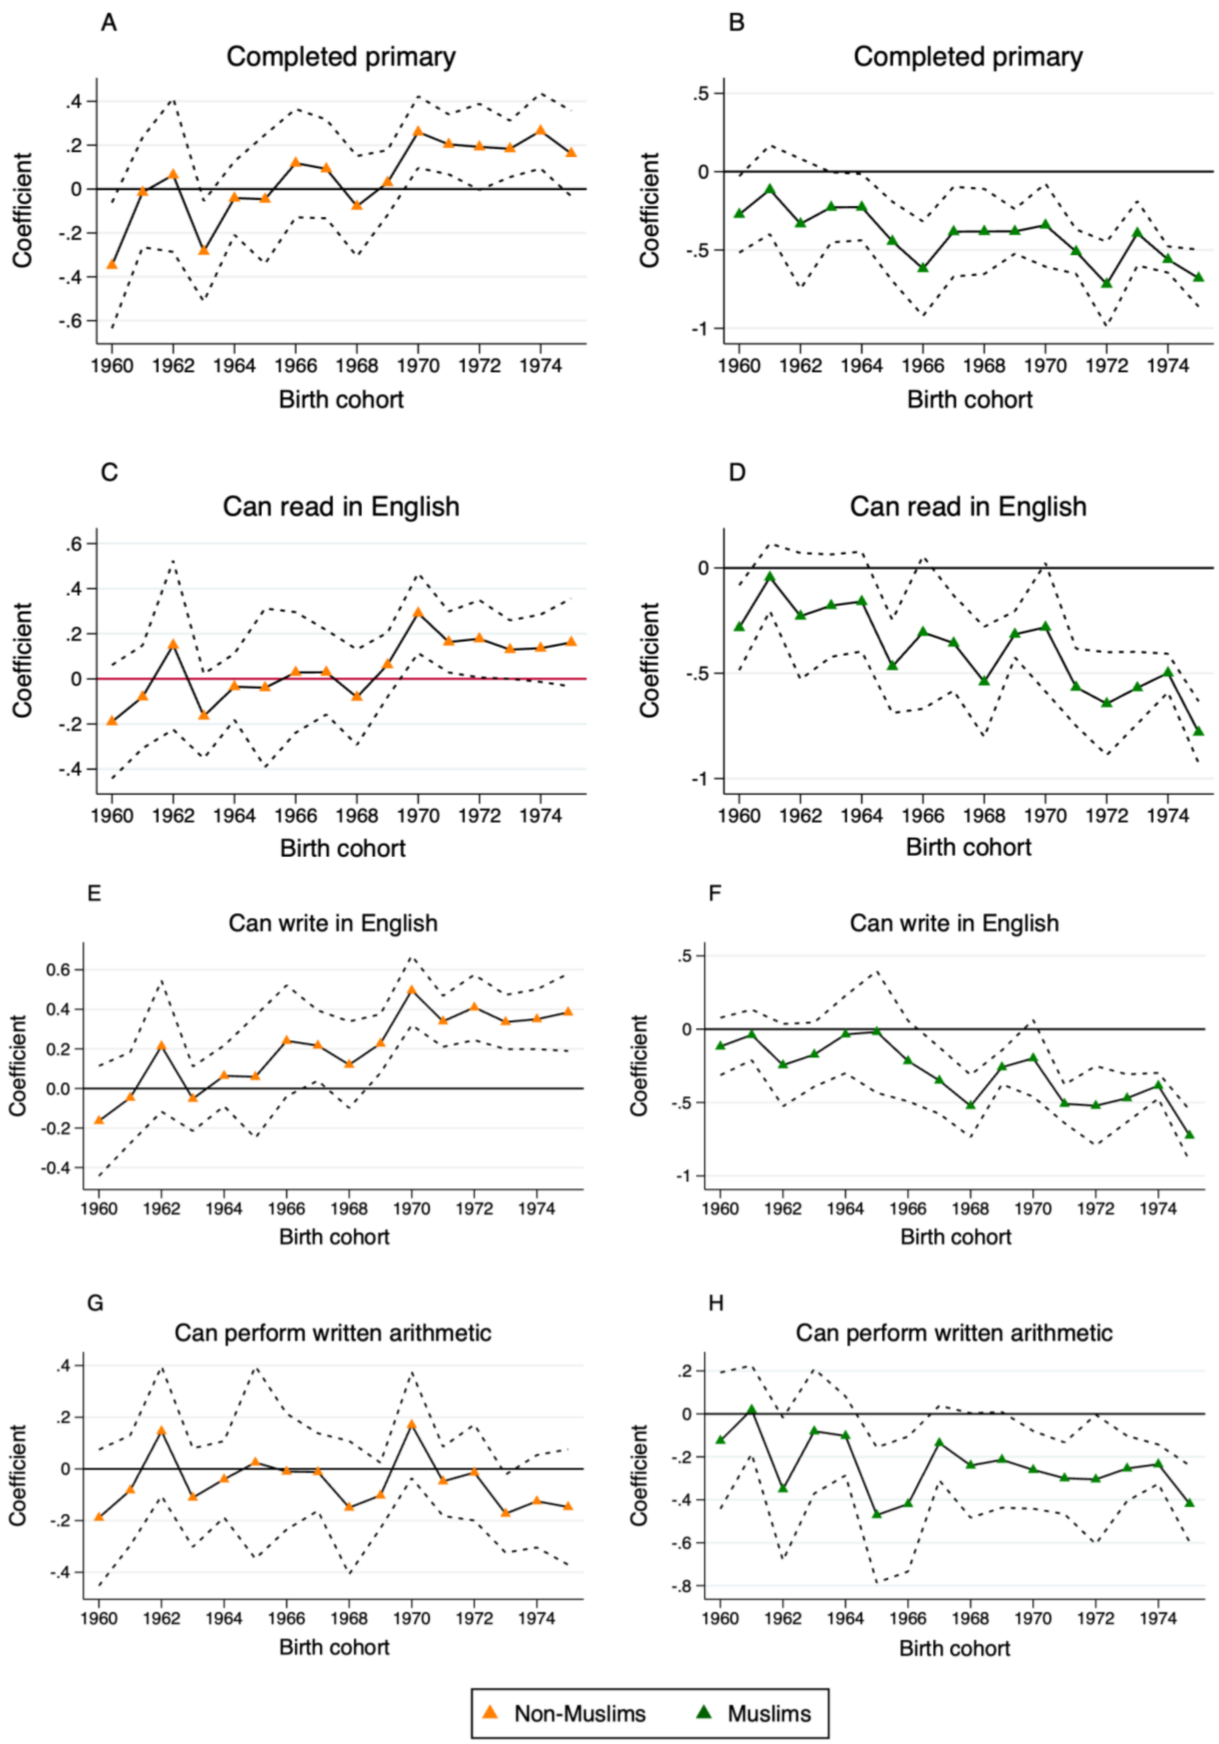

FIGURE 5. COEFFICIENTS OF THE INTERACTIONS BETWEEN YEAR OF BIRTH AND PROGRAMME INTENSITY IN THE STATE OF RESIDENCE WITH NINETY-FIFTH PERCENT CONFIDENCE INTERVAL BANDS

Notes.—The figure shows estimates of the effects of the programme on each birth cohort. Each panel represents a different regression. The title of each panel denotes the dependent variable. The omitted cohort (control cohort) in each regression is the 1959 cohort.

Selective migration.—I begin by exploring the extent to which selective migration could be driving my results. The HNLSS includes a migration module that contains the previous state of residence of respondents reporting not always living at their current address. For these individuals, I map the state-level data to their reported previous state of residence (instead of their current state of residence). This approach should increase the likelihood that I am using the state in which they attended school and reduce any bias in my estimates induced by selective migration.

Table 9 presents the results. Column (3), which presents estimates of the programme's effects on men, suggests that the paper's main finding that the programme increased male educational attainment and reduced male participation in agriculture still holds. The conclusion that the programme increased non-Muslim women's educational attainment at the primary level and improved learning outcomes also holds. Thus, it is unlikely that selective migration is driving these results.

Parallel trends.— I next test the plausibility of the parallel trends assumption. Following Brudevold-Newman (2021), I estimate over the pre-programme period (for the 1959-1964 cohorts) specifications with interactions between a trend variable and a dummy for high-intensity state, which equals one for states with above-median intensity[8]. Intuitively, if outcomes in low- and high-intensity states were on the same trend before the programme, estimates of the coefficients on these interactions should be statistically insignificant.

[8] These specifications are identical to specification (1) except that they do not include state specific linear trends and also control for interactions between cohort dummies and the 1976 unemployment rate in the state of residence

TABLE 9—ROBUSTNESS—USING REPORTED PREVIOUS STATE OF RESIDENCE

	Male	Female	
		Non-Muslim	Muslim
	UPE x Intensity	UPE x Intensity	UPE x Intensity
	(1)	(2)	(3)
Years of education	2.309***	3.435*	-1.397
	(0.755)	(1.782)	(2.665)
Completed primary	0.182**	0.378***	0.038
	(0.068)	(0.138)	(0.165)
Can read in English	0.076	0.349***	-0.130
	(0.077)	(0.116)	(0.159)
Can write in English	0.081	0.329***	-0.254
	(0.085)	(0.108)	(0.181)
Can perform written arithmetic	0.120	0.241**	-0.084
	(0.083)	(0.104)	(0.173)
Worked in agriculture	-0.163**		
	(0.060)		

Notes.—The table shows the estimated programme effects when assigning programme intensity based on respondents' previous state of residence. Each row represents a different specification. The dependent variables are in rows. All specifications control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, a rural-urban dummy, and interactions between cohort dummies and the 1976 unemployment rate in the current/previous state of residence. . Standard errors, reported in parentheses, are clustered at the state level, and all regressions have 37 clusters. The samples include men and women who are part of the 1959-1964 and 1970-1975 cohorts. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively.

TABLE 10—ROBUSTNESS—PARALLEL TRENDS ASSUMPTION

	Male	Female	
		Non-Muslim	Muslim
	Estimated coefficient on (High Intensity x trend)	Estimated coefficient on (High Intensity x trend)	Estimated Coefficient on (High Intensity x trend)
	(1)	(2)	(3)
Years of education	0.044	-0.169*	-0.243***
	(0.048)	(0.088)	(0.076)
Completed primary	0.006	-0.010	-0.033***
	(0.005)	(0.007)	(0.006)
Can read English	0.009	-0.007	-0.025***
	(0.006)	(0.007)	(0.007)
Can write in English	0.009	-0.006	-0.015**
	(0.006)	(0.006)	(0.007)
Can perform written arithmetic	0.005	0.003	-0.024***
	(0.006)	(0.007)	(0.007)
Worked in agriculture	0.012**	0.009	-0.001
	(0.005)	(0.008)	(0.006)

Notes.—The table shows the results of tests for the parallel trends assumption The number in each cell is the coefficient on an interaction between a dummy for high-intensity state and a trend variable. Each row represents a different specification. The dependent variables are in rows. All specifications control for cohort fixed effects, state fixed effects, dummies for religion, a rural-urban dummy, and interactions between cohort dummies and the 1976 unemployment in the state of residence. Standard errors, reported in parentheses, are clustered at the state level, and all regressions have 37 clusters. The samples include all pre-programme observations (1959-1964 cohorts). The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively.

TABLE 11—ROBUSTNESS—PLACEBO PROGRAMMES

	Male	Female	
		Non-Muslim	Muslim
	Estimated coefficient on (Placebo x Intensity)	Estimated coefficient on (Placebo x intensity)	Estimated coefficient on (Placebo x intensity)
	(1)	(2)	(3)
Panel A. Placebo 1987 x *intensity*			
Years of education	0.739	0.002	-1.293
	(1.031)	(1.211)	(1.275)
Completed primary	0.101	0.038	-0.016
	(0.113)	(0.082)	(0.075)
Can read English	0.044	0.002	-0.001
	(0.105)	(0.085)	(0.088)
Can write in English	0.068	0.165	0.034
	(0.092)	(0.123)	(0.097)
Can perform written arithmetic	0.019	0.047	0.028
	(0.090)	(0.084)	(0.092)
Worked in agriculture	0.037		
	(0.176)		
Panel B. Placebo 1988 x *intensity*			
Years of education	0.395	0.730	-2.098**
	(0.943)	(1.168)	(0.837)
Completed primary	0.043	-0.173*	-0.150***
	(0.095)	(0.099)	(0.054)
Can read English	-0.012	0.048	-0.156*
	(0.070)	(0.088)	(0.086)
Can write in English	0.067	0.173*	-0.142
	(0.064)	(0.101)	(0.087)
Can perform written arithmetic	0.023	0.154	-0.028
	(0.077)	(0.098)	(0.075)
Worked in agriculture	-0.138		
	(0.146)		
Panel C. Placebo 1989 x *intensity*			
Years of education	-1.313	1.301	-2.794***
	(1.175)	(1.052)	(0.750)
Completed primary	-0.080	0.237**	-0.236***
	(0.103)	(0.089)	(0.068)
Can read English	-0.112	0.108	-0.156**
	(0.087)	(0.107)	(0.068
Can write in English	-0.089	0.183	-0.120
	(0.084)	(0.118)	(0.080)
Can perform written arithmetic	-0.036	0.182	-0.094
	(0.084)	(0.112)	(0.072)
Worked in agriculture	-0.232*		
	(0.126)		
Panel D. Placebo 1990 x *intensity*			
Years of education	1.203	1.867	-2.169***
	(0.720)	(1.279)	(0.767)
Completed primary	0.001	0.195	-0.243***
	(0.040)	(0.116)	(0.065)
Can read English	0.062	0.192*	-0.197***
	(0.069)	(0.108)	(0.053)
Can write in English	0.027	0.162	-0.233***
	(0.077)	(0.115)	(0.066)
Can perform written arithmetic	0.091	0.237**	-0.211***
	(0.067)	(0.103)	(0.054)
Worked in agriculture	-0.156		
	(0.109)		

Notes.—The table shows the results of placebo programmes introduced in each of the years between 1987 and 1990. Each panel represents a different placebo programme. The number in each cell is the coefficient on an interaction between a continuous measure of programme intensity and a UPE indicator equal to one for individuals aged 2-6 years at the introduction of the falsified programme. Each row is a different specification. The dependent variables are in rows. All specifications control for cohort fixed effects, state fixed effects, state-specific linear trends, dummies for religion, a rural-urban dummy, and interactions between cohort dummies and the 1976 unemployment rate in the state of residence . Standard errors, reported in parentheses, are clustered at the state level, and all regressions have 37 clusters. For other educational and labour market outcomes, see table A5 in Appendix. The symbols *, **, and *** indicate significance at the levels of 10%, 5%, and 1%, respectively.

Table 10 presents the results. For males, all the estimates except the estimated effect on participation in agriculture are statistically insignificant (Column (1)). These estimates indicate that, before the programme, male engagement in agriculture in high-intensity states was increasing relative to that in low-intensity states. For non-Muslim women (column (2)) all the estimated coefficients are insignificant statistically, except those on junior secondary school completion and participation in skilled work. The coefficients for Muslims in column (3) are mostly negative and statistically significant, which is consistent with the finding from the cohort analysis that, in the pre-programme period, the educational outcomes of Muslim women in high-intensity states were declining relative to those of women in low-intensity states.

Placebo reforms.—Following Brudevold-Newman, I then examine the extent to which the results presented are statistical artefacts by implementing placebo tests examining hypothetical programmes introduced in each of the years between 1987 and 1990. I chose these years to be as close as possible to the UPE programme, while also avoiding confounding the effect of its ending in 1981 and the impact of Nigeria's other free education scheme introduced in 1999, the Universal Basic Education (UBE) programme. Mimicking the actual programme, I consider individuals who were aged 2-6 when the falsified programme was introduced as being part of treated cohorts, and those aged 12-17 form the control cohorts.

Table 11 presents the results. For men (column (1)), almost all the estimates are insignificant statistically, and the results for non-Muslim women in column (2) are also mostly insignificant statistically, with most estimates being smaller than those from the actual programme. However, for the 1989 and 1990 placebo reforms I find large and positive effects on female primary school completion and arithmetic skills. I cannot rule out the possibility that these are driven by partial exposure the UBE programme launched in 1999. This finding suggests that I cannot rule the possibility that some of my results are driven by unobserved time-varying factors.

The results for Muslim women, in contrast, are almost always negative and significant statistically. These results further support the finding from the cohort analysis that the negative effects among non-Muslim women from subsection B cannot be attributed to the UPE programme and might reflect unobserved factors.

VIII. Conclusion

The UPE programme increased educational attainment but does not seem to have improved long-run employment outcomes. The difference-in-difference strategy, exploiting cross-cohort and regional variation in exposure to the programme, suggests that the programme increased male educational attainment by one year, on average, and increased the probability of a man completing primary, junior secondary, and senior secondary schools by 13.7 percent, 21.7 percent, and 22.3 percent, respectively. The programme also increased the likelihood of a woman completing primary school by 36.8 percent and the probability of a woman reading in English by 28.6 percent. These positive effects on women seem to be concentrated among non-Muslims, possibly due to historical opposition from Muslim families towards `Western' education for girls around the time of the program's introduction (Clarke, 1978; Csapo, 1981).

The increase in men's educational attainment is consistent with the findings of Duflo (2001) and Khanna (2020), who conclude that the primary school construction programs in Indonesia and India increased male educational attainment by 0.12 - 0.19 years and 0.5 years, respectively. The larger effect of the UPE programme perhaps reflects the fact that, in contrast to these other two programs, it also abolished school fees. The positive impact on female primary school completion is consistent with the findings of Keats (2018) and Chicoine (2021), reporting increases in primary school completion among women who were exposed to free primary education programs in Uganda and Ethiopia, respectively.

The UPE programme does not seem to have increased the employability of men and women. The program reduced the probability of a man working in agriculture by around 11 percent, but it does not seem to have increased male participation in paid work, skilled work, and the public sector. The insignificant effect on female employability is consistent with the findings of Akresh (2023) but contrasts those of Chicoine (2021), who concludes that the free education program in Ethiopia raised female participation in skilled work.

Overall these findings suggest that a policy that abolishes school fees and build schools can generate significant gains in educational attainment. However, the insignificant effects on women's learning outcomes suggest that the increase in educational attainment may not always equate improved learning. My finding that the UPE programme increased secondary school completion only for men also suggests that policies that expand access to education may affect men and women differently.

The main limitation of this study is data quality. Most of the effects, especially on women, are imprecisely estimated, which is most likely due to noise in the data. Thus, I cannot rule out the possibility that some of the insignificant estimates are being driven by large standard errors rather than the absence of an effect per se.

The statistically insignificant effect of the programme on employability could be due to several factors. Identifying these factors would be interesting avenues for future research. For example, future research could explore the extent to which general equilibrium effects are driving these results. It would also be good to know whether the program had any effects on the earnings of men and women, which, due to data limitation, I do not consider.

X. References

Akresh, R., Halim, D., & Kleemans, M. (2023). Long-term and intergenerational effects of education: Evidence from school construction in Indonesia. The Economic Journal, 133(650), 582-612.

Asagwara, K. C. P. (1997). Quality of learning in Nigeria's universal primary education scheme—1976–1986. *The Urban Review*, *29*(3), 189-203.

Bass, V., & Nansamba, A. (2017). Information frictions in the labor market: Evidence from a field experiment in uganda.

Becker, G. S. (1962). Investment in human capital: A theoretical analysis. Journal of political economy, 70(5, Part 2), 9-49.

Becker, G.S. (1971). The Economics of Discrimination, second ed. The Univ. of Chicago Press, Chicago.

Brudevold-Newman, A. (2021). Expanding access to secondary education: Evidence from a fee reduction and capacity expansion policy in Kenya. Economics of Education Review, 83, 102127.

Chicoine, L. (2021). Free primary education, fertility, and women’s access to the labor market: Evidence from Ethiopia. The World Bank Economic Review, 35(2), 480-498.

Clarke, P. B. (1978). Islam, education and the developmental process in Nigeria. Comparative Education, 14(2), 133-141.

Csapo, M. (1983). Universal primary education in Nigeria: Its problems and implications. African Studies Review, 26(1), 91-106.

Duflo, E. (2001). Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment. American Economic Review, 91(4), 795-813.

General Household Survey. (2006). National Bureau of Statistics, Abuja.

Harmonized National Living Standard Survey. (2009). National Bureau of Statistics, Abuja.

Khanna, G. (2023). Large-scale education reform in general equilibrium: Regression discontinuity evidence from India. Journal of Political Economy, 131(2), 549-591.

Khatkhate, D. R. (1980). Capital scarcity and factor proportions in less developed countries. *Journal of Post Keynesian Economics*, *2*(3), 420-429.

Larreguy, H., & Marshall, J. (2017). The effect of education on civic and political engagement in nonconsolidated democracies: Evidence from Nigeria. Review of Economics and Statistics, 99(3), 387-401.

Lucas, A. M., & Mbiti, I. M. (2012). Access, sorting, and achievement: The short-run effects of free primary education in Kenya. American Economic Journal: Applied Economics, 4(4), 226-253.

Nandi, A., Haberland, N., & Ngo, T. D. (2023). The impact of primary schooling expansion on adult educational attainment, literacy, and health: Evidence from India's Sarva Shiksha Abhiyan. *International Journal of Educational Development*, *102*, 102871.

Osili, U. O., & Long, B. T. (2008). Does female schooling reduce fertility? Evidence from Nigeria. Journal of Development Economics, 87(1), 57-75.

Oyelere, R. U. (2010). Africa's education enigma? The Nigerian story. Journal of Development Economics, 91(1), 128-139.

Orazem, P. F., & King, E. M. (2007). Schooling in developing countries: The roles of supply, demand and government policy. *Handbook of development economics*, *4*, 3475-3559.

Mincer, J. (1958). Investment in human capital and personal income distribution. *Journal of political economy*, *66*(4), 281-302.

Nasidi, N. A., & Wali, R. M. (2023). The Challenges of Girl-Child Education in Ungogo Local Government Area of Kano State in Nigeria on 1999-2019. *Journal of Humanities and Social Sciences (JHASS)*, *5*(1), 27-36.