The Effect of Access to Primary School on the Timing of School Enrollment: Analysis of the Ethiopian Education Reform

Yared Seid[1,*], and Shiferaw Gurmu[2,*]

Article History: Received: 29 September 2022; Revised: 02 January 2024; Accepted: 18 January 2024

Abstract

This paper offers empirical evidence on whether access to primary school induces children to enroll in primary school at the legal enrollment age using household survey data from Ethiopia. We exploit the variation in the intensity of the impact of the education reform across districts in Ethiopia to identify the effect of access to school on the timing of enrollment. Using pre-reform enrollment rate in primary school to measure the variation in the intensity of the impact of the reform, we estimate difference-in-differences models. The results suggest that the reform has substantially increased the probability the child enrolls in grade 1 by age 7. It is also found out that the reform has decreased age at enrollment in grade 1 by about 4 months. These estimates highlight the important role access to school plays in inducing parents to enroll their kids in primary school at the legal enrollment age.

Keywords: school enrollment age, education, Ethiopia
JEL Classification: I25, O10, O12

[1] World Bank, Addis Ababa, Ethiopia; e-mail: yseid@worldbank.org
[2] Department of Economics, Andrew Young School of Policy Studies, Georgia State University, 55 Park Place NE, Atlanta, GA 30303, USA; e-mail: sgurmu@gsu.edu
* We thank Rachana Bhatt and Barry Hirsch for their invaluable suggestions. We have also benefited from the comments of Roy Bahl, James Marton, and Felix Rioja. In addition, we would like to thank participants at Georgia State University, Oxford University, and International Atlantic Economics Society Conferences. Any errors are the responsibility of the authors.

1. Introduction

One of the main features of the education system in developing countries is that the majority of students enroll in primary school long after the legal enrollment age, which is usually around 6 or 7 years. Barro and Lee (2001) find out that at least 50 percent of the students enrolled in grade 1 in 31 countries are older than the legal enrollment age. In Ethiopia, a country where the data for this study come from, the 2004 Welfare Monitoring Survey data show that more than 80 percent of children in rural areas enrolled in grade 1 after the legal enrollment age of 7. A number of other studies documented delayed primary school enrollment throughout the developing world (Bommier and Lambert, 2000; Glewwe and Jacoby, 1995; Wils, 2004; Moyi, 2010; Todd and Winters, 2011).

The standard human capital investment models fail to explain the widely observed delayed primary school enrollment as they predict that an individual invests in education in the early period of his/her life and reaps its benefits later in life. Besides, in communities where child labor is a common practice and most of the work children are expected to perform are physically demanding, it is optimal for parents to enroll the child as early as possible since the value of the child's time is lower when the child is younger. There is growing evidence that suggests delaying primary school enrollment is costly. In Ghana, for example, Glewwe and Jacoby (1995) calculated that delaying primary school enrollment by 2 years beyond 6 years, the legal enrollment age, costs an individual about 6 percent of his/her lifetime wealth. Also, children who enroll in school late have higher grade repetition and school dropout rates, and complete fewer years of schooling than those who enroll at the legal enrollment age (Wils, 2004). Given the high cost associated with delaying primary school enrollment, it is not well understood why most parents in developing countries enroll their children long after the prescribed age.

The bulk of the literature in this area focuses on the probability of enrollment, without considering age at enrollment. However, delayed enrollment cannot be tackled by general policies that are designed to increase enrollment rates since delayed enrollment is not confined to countries that have lower enrollment rates (Moyi, 2010; Lloyd and Blanc, 1996). Very few studies analyzed why students in developing countries delay primary school enrollment. Loosely speaking, the explanations these studies provided can be grouped into three: poor child health, liquidity constraint, and limited (or lack of) access to school.

Poor child health slows down the child's development process and renders the child less ready to attend school at the legal enrollment age. Hence, at legal enrollment age, say, a malnourished child would be too weak to be able to walk the (typically longer) distance to school (Bommier and Lambert, 2000; Partnership, 1999). Besides, poor child health lowers the learning ability of the child and thus it is optimal to delay enrollment until the negative effect of poor child health on mental readiness decreases after a few years when the child gets older (Glewwe and Jacoby, 1995). A liquidity constraint explanation, on the other hand, suggests that resource constrained families might need to employ the child in family activities until the family accumulates sufficient saving to finance the child's schooling (Jacoby, 1994). Finally, if there is limited access to schooling, school officials may ration enrollment in primary school, and the rationing tends to favor older children who are typically on the waiting line for a relatively longer time (Bommier and Lambert, 2000). Note that shortage of schools means schools are widely dispersed and we expect children to walk for a relatively longer distance. Since malnourished children are too weak to walk for longer distance to school, school shortage may interact with child health and have differential impact across children on the health distribution.

Most families in developing countries, particularly those in rural areas, do not have access to primary schools. In recent years, however, many developing countries have made primary schools more accessible. There is evidence that suggests making schools more accessible has increased primary school enrollment, but we know little about the effect of access to school on the timing of enrollment. This paper, thus, attempts to bridge this gap in the literature by offering empirical evidence on the effect of access to school on the timing of primary school enrollment using a household survey data from Ethiopia.

One of the empirical challenges of assessing the effect of access to school on the timing of enrollment is endogeneity of access to school; that is, families that live closer to school may be inherently different and their children may enroll in school on time regardless of their proximity to school. In situations like these, most researchers attempt to mitigate the bias by either finding appropriate instrumental variable or using a "natural experiment" that affects the endogenous variable but not the outcome variable. Some prior studies exploit government programs as exogenous source of variation in economic variables. For example, Todd and Winters (2011) and McEwan (2013), respectively, exploit the government programs in Mexico (called Oportunidades) and Chile as exogenous source of variation in child health to investigate the effect of child health on the timing of school enrollment. This study employs a similar approach and uses an education policy shock that happened in Ethiopia between the mid-1990s and mid-2000s as exogenous source of variation in access to primary school.

The Ethiopian government has launched a series of five-year Education Sector Development Programs (ESDPs) since 1997 with a prime objective of achieving universal primary education by 2015. To date, 6 five-year ESDPs have been implemented. During the first two ESDPs that covered 8 academic years between 1997/98 and 2004/05, 2,398 new primary schools were built (World Bank, 2005). The program has substantially decreased distance to primary school at a national level from its average of 2.73 Km in 1996 to that of 1.25 Km in 2004. Though such a large number of primary schools were built in a short period of time and the program has substantially decreased distance to primary school at a national level, the decrease in distance to primary school vary widely across states and zones.[23] For example, distance to primary school has decreased by 100 percent in East Wellega zone while the decrease was only 2.81 percent in South Gondar zone during the same period.

We exploit the variation in the intensity of the impact of the program across states to identify the causal effect of access to primary school on the probability of enrollment in grade 1 by age 7, the legal enrollment age. Narrowing down education inequalities across states by building more schools in rural and under-served areas was at the core of the program's objective. In fact, the program explicitly targeted increasing primary school enrollment from its 30 percent national average at the beginning of the program to at least 50 percent by the end of the program. Thus, we should expect more schools to be built in areas that had lower primary school enrollment rate in the pre-program period. Accordingly, states that had pre-program primary school enrollment rate below 30 percent are assigned into treatment group, whereas those states above 30 percent enrollment rate are assigned into control group. Then, difference-in-differences models are estimated where the dependent variable is a binary indicator for enrollment in grade 1 by age 7. To estimate the models, we use household survey data - called Welfare Monitoring Survey (WMS) data - administered by Ethiopia's Central Statistical Agency during the periods 1996 and 2004. The main advantage of using datasets from these survey rounds is that they have information on important variables just before the beginning of the program (i.e., 1996) and around the end of the program (i.e., 2004). Moreover, this is a period where the Ethiopian government has started mass building primary schools across the country, focusing primarily on under-served communities.

[23] Ethiopia is a federal country with three levels of governments: federal, state (or regional), and local governments. Zones are mid-level administrative unit of the government that are equivalent to US counties.

The results from the difference-in-differences models reveal that the education program has increased the probability the child enrolls in grade 1 by age 7 by more than 35 percentage points. The results also suggest that the reform has decreased age at enrollment in grade 1 by about 4 months. These estimates highlight an important role access to school plays in inducing parents to enroll their kids in primary school at the legal enrollment age.

The remainder of the paper is organized as follows. The following section briefly reviews the literature, and Section 3 describes the education reform in Ethiopia. Section 4 explains the data and presents descriptive statistics. The impact of the education program on both access to school and primary school enrollment is discussed in Section 5. Section 6 presents the evidence on the impact of the education program on the timing of enrollment. While doing so, Section 6 discusses the identification strategy and presents the econometric results. The final section concludes.

2. Literature Review

Delayed primary school enrollment is observed in a number of developing countries. For instance, it has been documented in most Sub-Saharan African countries (Barro and Lee, 2001), in Tanzania (Bommier and Lambert, 2000), Ghana (Glewwe and Jacoby, 1995; Seshie-Nasser and Oduro, 2016)), Mozambique (Wils, 2004), Malawi (Moyi, 2010), Mexico (Todd and Winters, 2011), and rural China (Chen, 2015). Contrary to the fact that delayed enrollment is common in most developing countries, we have limited understanding of its causes.

Prior studies on the topic suggest various explanations on why parents delay their children's enrollment in primary school. First, malnutrition could cause delayed primary school enrollment. This is because malnutrition lowers children's learning ability and hence it is optimal to delay enrollment until the negative effect of malnutrition decreases after a few years when the child gets older (Glewwe and Jacoby, 1995). Using a policy intervention that improved child health in Mexico as exogenous source of variation in child health, Todd and Winters (2011) find out that early health and nutrition intervention has increased the probability a child enrolls on time in primary school. On the contrary, McEwan (2013) finds out that a similar policy intervention that made higher calorie meal available to vulnerable children in Chile has no effect on enrollment in grade 1 at the legal enrollment age. The author attributes this to the lower incidence of child malnutrition in Chile, and high proportion of children in Chile being enrolled in school on time. On the other hand, in Ghana and Tanzania, Partnership (1999) found out that malnutrition, measured by height-for-age, delays enrollment in primary school.

Second, De Vreyer et al. (1999) model a household behavior, where households diversify their investment among three assets: physical assets, general human capital acquired through schooling, and specific human capital acquired through child labor. If the return to specific human capital at younger age is higher than that of general human capital, then parents do not send their children to school at the legal school enrollment age. Third, delayed school enrollment could be the result of liquidity constraints. When households are resource constrained, a child might need to be employed in family activities until the family accumulates sufficient saving to finance the child's schooling (Jacoby, 1994).

Finally, delayed school enrollment could be the result of supply side problems. If there is shortage of school, school officials may ration enrollment in primary school, and the rationing tends to favor older children who are typically on the waiting line for a relatively longer time. On the other hand, shortage of school may mean students have to walk longer distance to school. In this case, delayed enrollment could be due to the fact that children may not be mature enough to walk the distance to school at the legal enrollment age (Bommier and Lambert, 2000). In societies where there is high incidence of child malnutrition, shortage of schools exacerbates the problem of delayed enrollment since developmentally stunted children take relatively longer time to be physically strong and be able to walk the longer distance to school. On the other hand, walking longer distance to school increases the propensity a child walks through unsafe neighborhoods. Thus, parents that are concerned about the safety of their children may refrain from sending their children, especially their daughters, to school at the legal enrollment age.

Though access to school is one of the most important factors that determine the timing of enrollment, identifying its effect on the timing of enrollment is complicated by the relationship between school proximity, socioeconomic status, parental taste for education, and other characteristics that affect the timing of enrollment. For instance, being economically disadvantaged is correlated with poor taste for education and living further away from schools. All these factors affect the timing of enrollment, but they cannot be perfectly controlled in the regression framework. A credible identification of the effect of access to school on the timing of enrollment, thus, requires exogenous source of variation in access to school that does not affect the timing of enrollment.

In situations like these, government programs can be used as exogenous source of variation in the independent variable. For example, to test the hypothesis that malnutrition delays school enrollment, Todd and Winters (2011) used the government program in Mexico called Oportunidades as exogenous source of variation in child health. McEwan (2013) also used a similar intervention in Chile as exogenous source of variation in the amount of calorie intake among children to identify the causal effect of child health on the timing of enrollment. We follow a similar approach and use the education reform that happened in Ethiopia between 1996 and 2004 as exogenous source of variation in proximity to school to identify the causal effect of access to primary school on the timing of enrollment.

3. The Education Reform in Ethiopia

Following the change in government in May 1991, Ethiopia has undergone a number of policy changes in almost all sectors of the economy. The education sector is one of the sectors that has gained the attention of the government since then. Consequently, it has undergone many policy changes and received a large and increasing budget share of the government. Among the many changes the sector experienced recently, the implementation of a series of five-year Education Sector Development Programs (ESDPs) is the major one. We exploit the variation in the intensity of the impact of the education program across districts to identify the causal effect of access to primary school on the timing of enrollment in grade 1.

The ESDPs started in 1997 with the objective of achieving universal primary education by 2015. Reducing educational inequalities by increasing access to primary school, mainly in rural and under-served areas, was at the core of the ESDPs. To date, 6 five-year ESDPs have been implemented. We will focus on the first two ESDPs in this paper as their duration aligns with the survey years of the data used in this paper and this is the period where Ethiopia has begun building a large number of schools to substantially increase access to primary education, albeit from a very low base.

The first ESDP covered five academic years between 1997/98 and 2001/02. Over the five years period of the first ESDP, it was planned to build 2,423 new primary schools, to upgrade 1,814 primary schools, and to renovate 1,220 primary schools in order to accommodate 3.9 million additional students (World Bank, 1998). The expected outcomes were substantial increase in access to primary school especially in rural areas where the majority of newly built schools were to be located. Moreover, it was expected to increase gross primary school enrollment rate from its 30 percent level by the beginning of the first ESDP to 50 percent by the end of the first ESDP.

The second ESDP also covered five academic years between 2000/01 to 2004/05. The first two years of the second ESDP overlapped with the last two years of the first ESDP. Thus, in effect, the second ESDP had covered three unique academic years between 2002/03 and 2004/05. The reason for the overlap in the duration of the first and second ESDPs is to align the second and consecutive ESDPs with the political election cycle and the five-year term of the elected government in office. Though it was planned to build 2,423 primary schools during the first ESDP alone, a total of 2,398 new primary schools were built during the first two ESDPs, and, in line with the focus of the program, 86 percent of the new schools were built in rural areas (World Bank, 2005).

As the first two ESDPs covered 8 academic years between 1997/98 and 2004/05, household survey data collected in 1996 and 2004 are used in this paper so that the 1996 and 2004 data are, respectively, used as pre and post program data. The next section discusses the data used in this paper and presents descriptive statistics.

4. Data

The analysis in this paper is based on the Ethiopian Welfare Monitoring Survey (WMS) data, which was administered by the Ethiopia's Central Statistical Agency in 1996 and 2004. The WMS is a cluster-based nationally representative repeated cross section household survey. The 1996 and 2004 WMS covered 11,569 and 36,303 households, respectively, and the surveys contain a wide range of information on household demographics, household assets, availability and use of different facilities (including schools), and other important economic variables.

For each household member aged five and above, we observe whether an individual was attending school during the survey years and a year prior to the survey years. We also observe the grade in which an individual was registered in these two consecutive years. Using this information, we restricted the sample to first time grade 1 enrollee in the two survey years. Since grade repetition is common in Ethiopia as it is in most developing countries, it is important to mention that one of the advantages of these data is that they allow us to observe first time grade 1 enrollees. Hence, bias from measurement error of age at enrollment, which can be caused by grade repetition, is not a serious concern here.

Table 1 presents descriptive statistics for a sample of 2,368 children used in the econometric analysis. Detailed definition of variables is given in Table A.1 of the Appendix. The table shows that children in rural areas, on average, enroll in grade 1 at least 2.5 years after the legal enrollment age of 7. The extent of delayed enrollment in rural area is also reflected by the small proportion of children that were enrolled in grade 1 by age 7, which was 11 percent in 1996 and 18 percent in 2004. Although delayed enrollment in urban areas is less common as compared to rural areas, a non-trivial number of children in urban areas enroll in grade 1 after the legal enrollment age. For instance, Table 1 depicts that only about 50 percent and 53 percent of children in urban areas were enrolled in grade 1 by age 7 in 1996 and 2004, respectively. Moreover, children in urban areas delay enrollment in grade 1 by about a year in 1996 and 10 months in 2004. To summarize, a sizable proportion of children enroll in grade 1 few years after the legal enrollment age of 7 years. However, children in rural areas are more likely to delay enrollment, and when they do, they delay enrollment by more years than their urban counterparts.

The enrollment rate for girls has been disproportionately lower than for boys for a long time, particularly in rural Ethiopia. Table 1, however, shows that the proportion of girls enrolled in grade 1 has been increasing during the period of analysis, both in rural and urban areas. Given narrowing down gender gap in primary school enrollment was one of the objectives of the program, it is interesting to see increasing proportion of girls were enrolled in grade 1 during this period.

Generally, parents in Ethiopia are less educated, with the highest average years of schooling being 5 years for fathers and 3 years for mothers. As expected, parents in urban areas are more educated than their rural counterparts. Parental years of schooling has slightly increased in rural areas between 1996 and 2004. Though it is not clear why this is the case, it could be partly because of ongoing adult education in Ethiopia.

Table 1: Descriptive statistics for a sample of children who were enrolled in grade 1 by year and location

1996 2004	Rural	Urban	Rural	Urban
Enrolled in grade 1 by age 7 (yes=1)	0.111	0.498	0.176	0.528
	(0.314)	(0.501)	(0.381)	(0.500)
Age at enrollment	10.203	7.967	9.635	7.752
	(2.128)	(1.760)	(2.151)	(1.587)
Girl (yes=1)	0.304	0.474	0.452	0.534
	(0.461)	(0.501)	(0.498)	(0.500)
Birth order	2.538	3.414	2.509	2.859
	(1.334)	(1.936)	(1.414)	(1.646)
Household size	7.184	7.395	7.032	6.662
	(1.895)	(2.409)	(1.896)	(2.110)
Dad’s years of schooling	1.108	5.107	1.629	5.041
	(2.165)	(3.687)	(2.580)	(3.679)
Mom’s years of schooling	0.149	3.386	0.570	3.248
	(0.768)	(3.642)	(1.635)	(3.756)
Dad’s age	45.364	45.558	44.803	43.931
	(9.725)	(11.411)	(10.615)	(10.798)
Mom’s age	37.038	36.293	35.867	34.614
	(7.840)	(7.353)	(7.809)	(7.710)
HH has piped water (yes=1)	0.044	0.716	0.151	0.783
	(0.206)	(0.452)	(0.358)	(0.413)
HH has electricity (yes=1)	0.013	0.847	0.020	0.707
	(0.11	(0.361)	(0.140)	(0.456)
HH has pit latrine (yes=1)	0.092	0.758	0.253	0.741
	(0.289)	(0.439)		
HH owns land (yes=1)	0.997	0.703		
	(0.056)	(0.458)		
HH owns farm animal (yes=1)	0.633	0.655		
	(0.483)	(0.476)		
Proportion of hhs with piped water	0.035	0.819		
	(0.137)	(0.301)		
Proportion of hhs with electricity	0.016	0.677		
	(0.087)	(0.361)		
Proportion of hhs with pit latrine	0.083	0.705		
	(0.204)	(0.267)		
Proportion of hhs with land	0.974	0.547		
	(0.048)	(0.269)		
Proportion of hhs with farm animal	0.493	0.515		
	(0.300)	(0.295)		
State-level unemployment rate*	2.479	5.109		
	(4.743)	(6.068)		
Observations	316	290		

Standard deviations are in parentheses.
*Source: The 1994 and 2007 Ethiopian Census.
Proportion of households is defined over the locality of the child’s residence which is roughly equivalent to a village or an urban neighborhood.

Household assets and amenities variables depicted in Table 1 show that families in rural areas have fewer household assets and live in poor housing conditions compared to those in urban areas. However, household assets and housing condition have improved during the period of analysis for households in both rural and urban areas. To control for the economic condition of the locality of the child's residence, we control for the proportion of households that owns different types of household assets and amenities in the locality of the child's residence; that is, enumeration area which is used as primary sampling unit in the survey design and is roughly equivalent to a village or an urban neighborhood. As expected, the table shows that rural localities are relatively poorer than their urban counterparts. Finally, Table 1 also shows that unemployment rate varies by location of residence, with urban unemployment rate higher than rural unemployment rate.

5. The Impact of the Education Program on Access to School and Primary School Enrollment

The education program substantially increased access to school in Ethiopia. As mentioned earlier, 2,398 new primary schools were built over a period of 8 years as a result of the program (World Bank, 2005). Besides, data from the 1996 and 2004 Ethiopian Welfare Monitoring Survey show that the average distance to primary school had decreased at a national level by 1.48 kilometers (Km) between 1996 and 2004, which is more than a 50 percent decrease from its average of 2.73 Km in 1996 to that of 1.25 Km in 2004.

Although the program has substantially decreased distance at the national level, Figure 1 shows that the change in distance to primary school during this period vary widely across zones. Of the total 52 zones surveyed both in 1996 and 2004, distance to primary school decreased in 43 zones, ranging from a 100 percent decrease to that of 2.81 percent decrease. On the other hand, distance to primary school increased in 9 zones during the same period, ranging from a 1.13 percent increase to that of 203 percent increase.

Figure 1: Percentage change in distance to primary school between 1996 and 2004 by zones in Ethiopia

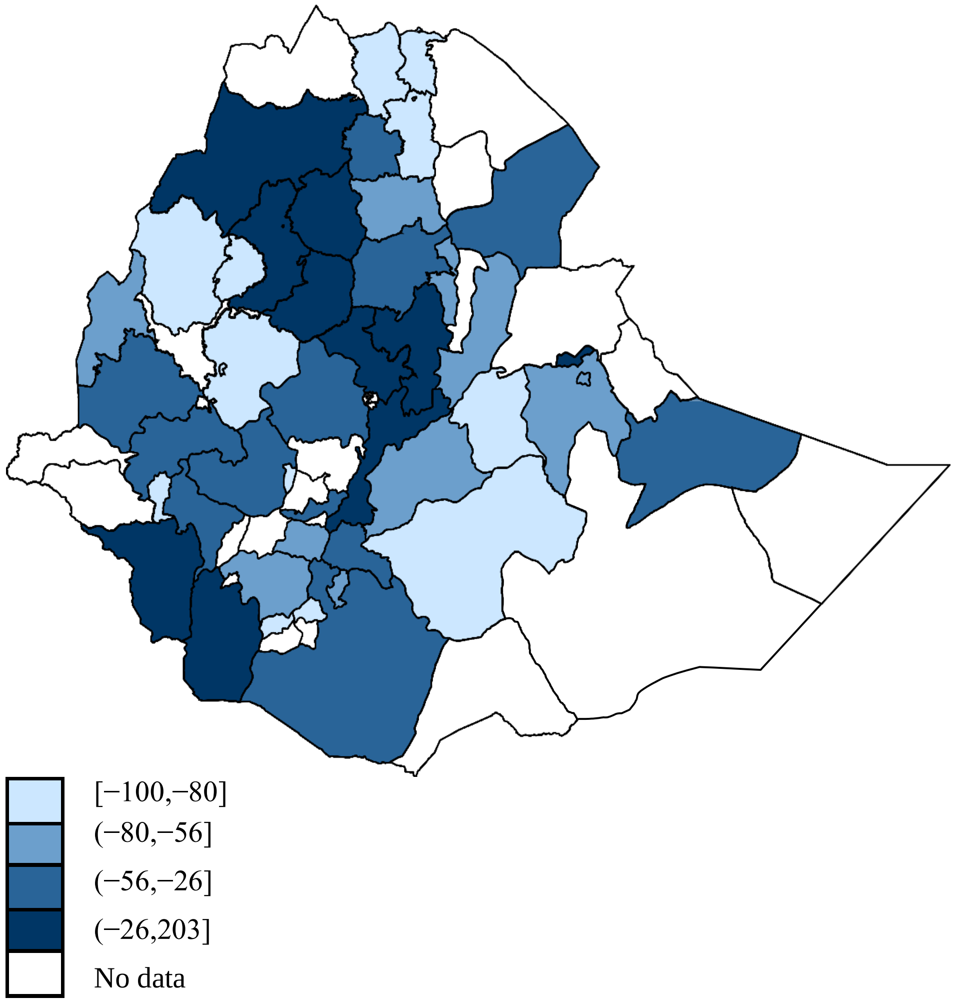

Similarly, enrollment in primary school has increased substantially in recent years. Figure 2 depicts the trend in enrollment rate in primary school in the last three decades using data from the World Bank. For the most part of the 1980s, enrollment rate was stable around 40 percent, except in the late 1980s where it started to decline. The decline is mainly because of the aggravated civil war between the military government in power at that time and the rebellion group that finally threw the military government out of power in 1991.

Starting the early 1990s, enrollment rate has started to increase and reached its 1980s level around 1997. Enrollment rate has been continuously increasing since then. The increase in enrollment rate during the period of analysis (which is marked between the two vertical lines in Figure 2) is attributed to the education program that has been in place. Even if the focus of this paper is on the education program that was implemented between 1996 and 2004 (more specifically, the first and second Education Sector Development Programs), the next phase of the program (the third Education Sector Development Program) has been implemented by the end of the second phase of the program. Therefore, we should not expect the growth in enrollment rate to decrease or plateau after 2004. That is why the curve in Figure 2 continuously increases even after 2004.

Figure 2: Primary School Enrollment Rate Trend in Ethiopia

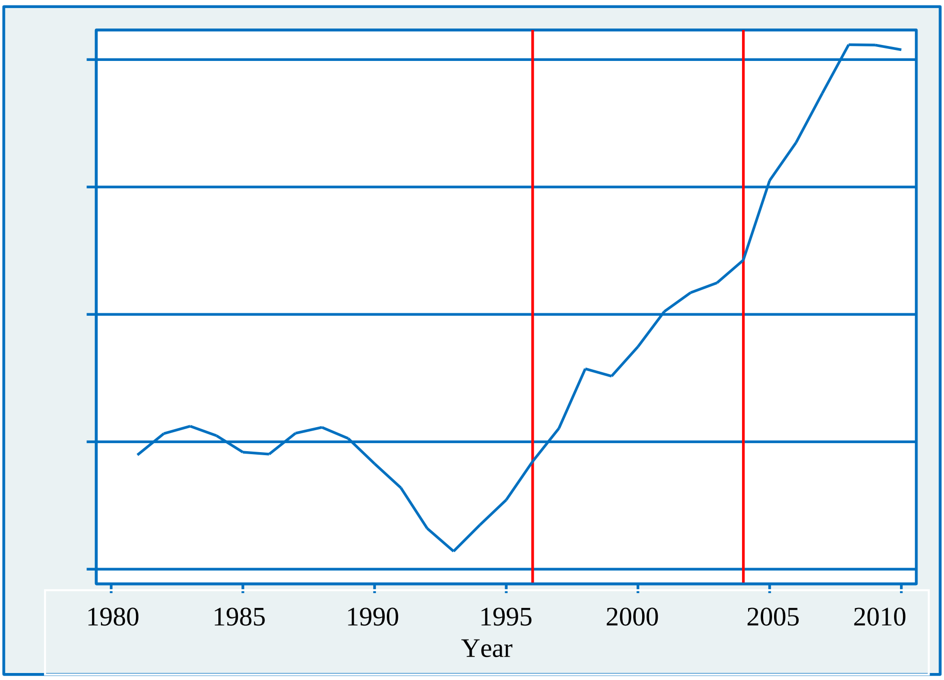

Source: World Bank

One feature of the education program is narrowing down educational inequalities across states and between rural and urban residents. This is reflected in the allocation of the newly built schools where 86 percent of them were built in rural areas (World Bank, 2005). The program also explicitly targeted increasing primary school enrollment from its 30 percent national average at the beginning of the program to at least 50 percent by the end of the program. We should, therefore, expect more schools to be built in states that had less than 30 percent enrollment rate before the program. Accordingly, we assign states with less than 30 percent enrollment rate before the program in the treated group and those above 30 percent in the control group.

Using the 1994 Ethiopian census data, Table 2 presents the primary school enrollment rate before the program by state and treatment status. Three states had enrollment rate above 30 percent before the program. These three states are assigned into a control group and all the remaining states are assigned to a treatment group. It is crucial to mention that the three states in the control group are largely urban in nature and lead primarily non-agrarian economies. But note that all the other states also have major urban areas[24] although the majority of their residents live in rural areas. Given the program focused on building the majority of the schools in rural areas, it is expected states in the control group to be predominantly urban in nature.

Table 2: Enrollment Rate in Primary School (Grades 1-8) During the Year before the Education Program

State/Region	Enrollment Rate	Treated State?
Tigray	15.2	Yes
Afar	2.96	Yes
Amhara	7.64	Yes
Oromiya	9.52	Yes
Somali	2.03	Yes
Benishangul Gumuz	9.94	Yes
SNNP	10.9	Yes
Harari	31	No
Addis Ababa	62	No
Dire Dawa	31.6	No

Source: The 1994 Ethiopian Census

[24] Central Statistical Agency of Ethiopia defines two types of urban areas: *major urban areas* and *other urban areas*. This classification depends on the nature of economic activity and the number of residents. All state capitals are considered as major urban areas, and they are typically more developed and have larger population size relative to other urban areas.

6. Estimation Strategy and Identification

6.1. Econometric Method

Let us denote enrollment status of child *i* residing in state's in year *t* by $enroll_{ist}$. Taking into account individual differences in observable characteristics, the probability that the child enrolls on time it can be written in a generic form

$$Pr(enroll_{ist} = 1|d_{it}, \mathbf{X}_{ist}) = G(\gamma d_{it} + \mathbf{X}_{ist}\boldsymbol{\beta}), \qquad \textbf{(1)}$$

where d denotes distance to primary school, X_{ist} represents a vector of explanatory variables, and G is a known function of covariates.

If access to primary school (d_{it}) is endogenous in equation (1), estimating equation (1) provides biased estimate of γ and hence it cannot be interpreted as the causal effect of access to school on the probability of enrollment on time. There are a number of reasons why we expect access to primary school to be endogenous in equation (1), including unobserved parental taste for education. Generally, families that live closer to schools may be inherently different and their children may enroll in school on time regardless of their proximity to school. If there is exogenous source of variation in proximity to school that does not affect the outcome variable, the causal effect of access to school on the timing of enrollment can be identified. We exploit the variation in the intensity of the impact of the education program across states in Ethiopia to identify the causal effect of access to school on the timing of enrollment.

Difference-in-Differences Approach

Ideally, we would compare the probability of enrollment in grade 1 by age 7 ($enroll_i$) for the same set of children when they are exposed to the education program ($enroll_i$ |education program) and when they are not ($enroll_i$ |no education program). In this ideal case, the average treatment effect would be the differences in the expected values under the two scenarios.

However, the same set of children cannot be observed under both scenarios since the child is either exposed to the program or not. Hence, to estimate the average treatment effect, data on two groups of randomly assigned children where one group is exposed to the program (treatment group) while the other is not exposed to the program (control group) are required. As long as assignment of children to treatment

(Treated = 1) and control (Treated = 0) groups are random, the average treatment effect can be obtained by first difference model.

If, however, children in the two groups differ initially and have different timing of enrollment in the absence of the program, we have to control for the pre-existing difference between the two groups. If we have information on observations both before the education program occurred (After = 0) and after the program occurred (After = 1), then a difference in differences approach can be used to separate the pre-existing difference from that of the treatment effect. Specifically, we can estimate:

$$P\ r(enroll_{it} = 1) = G(\alpha_0 + \eta_0\ Treated_{it} + \tau_0 After_{it} + \gamma_0\ Treated_{it} * After_{it}) \quad \textbf{(2)}$$

In linear probability model, η0 in equation (2) estimates the pre-existing difference between children in the two groups, τ0 estimates the change in the outcome that occurred over time due to other factors, and γ0 estimates the impact of the education program. Estimating γ0 in equation (2) assumes children in the two groups would experience the same time trend (τ0) in the absence of the program, so that once initial difference (η0) and time trend are controlled for, the remaining difference between children in the treatment and control groups can be attributed to the program.

As mentioned earlier, state level pre-program enrollment rate in primary school is used to group states (and hence students) into treatment and control groups. Specifically, students that live in states that had pre-program primary school enrollment rate below 30 percent are assigned into treatment group, whereas students that live in states with pre-program primary school enrollment rate above 30 percent are assigned into control group. The argument is that relatively more schools should be built in areas where the pre-program enrollment rate in primary school is lower since the program explicitly targeted narrowing down education inequalities across states by building more primary schools in areas where primary school enrollment rate was lower before the education program. Hence, if proximity to primary school induces children to enroll on time, in the post-program period, we expect to see children in the treated states to be more likely to enroll in primary school on time relative to those that live in control states.

The basic identification strategy can easily be demonstrated by a simple difference- in-differences table. Table 3 presents the difference in differences in age at enrollment in grade 1 between children in the treated and control states before and after the education program. The first column of Table 3 displays that, before the program, children in the treated group enrolled in grade 1 at age 9.5 on average while those in the control group enrolled at mean age of 8.0, a difference of 1.6 years. The difference, however, narrowed down to 0.6 years after the program. Thus, the difference in the differences in mean age at enrollment in grade 1 is about −0.992 years with a standard error of 0.339. This indicates that the unadjusted treatment effect is about -1 year, implying that the educational reform program has reduced age at enrollment in first grade by age 7 by about 1 year.[25]

Table 3: Age at Enrollment by Treatment Group Before and After the Program (Observations 2368)

Group	Before the Change	After the Change	Time Difference
Treated	9.546*** (0.103)	9.372*** (0.053)	-0.174 (0.116)
Untreated	7.976*** (0.238)	8.794*** (0.211)	0.818*** (0.318)
Group Difference	1.570*** (0.260)	0.577*** (0.218)	-0.992*** (0.339)

Notes: Standard errors are given in parentheses. The standard error associated with the treatment effect (highlighted) is clustered by enumeration area, the primary sampling unit. Statistical significance: * p <0.10, ** p < 0.05, and *** p < 0.01

The difference in differences can be interpreted as the causal effect of the program under the assumption that in the absence of the program the decrease in age at enrollment would not have been systematically different in treated and control states. If this assumption is not satisfied, the difference in differences presented here cannot be interpreted as the "true" treatment effect. In the paragraphs below, we present a difference-in-differences model that adjusts for observable differences between individuals in the treated and control groups in the regression framework.

Using observations sampled from 10 states in Ethiopia and controlling for individual, household, and community-level characteristics; state fixed effects; and state-by-year fixed effects to improve precision, we estimate:

$$Pr(enroll_{ist} = 1) = G(\alpha + \eta Treated_s + \tau After_{it} + \gamma Treated_s * After_{it} + \boldsymbol{\beta}_1\mathbf{X}_{it} + \boldsymbol{\beta}_2\mathbf{W}_{ht} + \boldsymbol{\beta}_3\mathbf{C}_t + \boldsymbol{\beta}_4\mathbf{S} + \boldsymbol{\beta}_5\mathbf{S} * \mathbf{Y}) \quad \textbf{(3)}$$

[25] A counterpart of Table 3 which uses means of enrollment dummy is presented in Table A.2 in Appendix A. Table A.2 shows that the unadjusted treatment effect is 0.107, suggesting the program has increased the probability of enrollment in grade 1 by age 7 by 10.7 percentage points. Note that the same result can be obtained from OLS regression of equation (2).

where $enroll_{ist}$ is a dummy variable which takes a value of 1 if child i in state s in year t is enrolled in grade 1 by age 7; $Treated_s$ is a binary indicator for states that had pre-program primary school enrollment rate below 30 percent; $After_{it}$ is a dummy variable equal to 1 if the child is being observed after the program, and zero otherwise; $\mathbf{X}_{it}$, $\mathbf{W}_{ht}$, and $\mathbf{C}_t$ are vectors of individual, household, and community level characteristics, respectively; **S** is a vector of state dummies to control for (time invariant) state fixed effect; and ($\mathbf{S} * \mathbf{Y}$) is a vector of binary indicators for the interaction of state and year dummies to control for state-specific shocks over this period which are correlated with the education program.[26]

The primary (explanatory) variable of interest is the interaction term, $Treated_s *$

$After_{it}$, and γ captures the treatment effect, i.e., the effect on the probability a child enrolls in grade 1 by age 7 due to the child lives in the treated states (relative to those that live in the control states) after the program has occurred. While estimating equation (3), the standard errors are clustered by enumeration area, a primary sampling unit, to account for correlation in the error terms within enumeration area over time. For the most part, we assume **G** is standard normal cumulative distribution function and estimate a probit model, in which case the estimate of the average treatment effect and its standard error are computed as suggested by Puhani (2012).

6.2. Econometric Results

Table 4 presents Probit estimates of the average treatment effects (effect of education reform on on-time primary school enrollment) for various specifications depending on the control groups used in equation (3). The dependent variable is a binary indicator for enrollment in grade 1 by age 7. All specifications include state and state-by-year fixed effects to the extent possible due to potential collinearity between these effects and post-treatment and treatment dummy variables. Table A3 in the appendix provides Linear Probability Model (LPM) treatment effect estimates

[26] A slightly different version of the model presented in equation (3) is the one that replaces the dummy variable for treated group, Treated_S, by a continuous pre-program state level primary school enrollment rate variable, EnrolRates, i.e.,

$$Pr(enroll_{ist} = 1) = G(\alpha + \eta EnrolRates_S + \tau After_{it} + \gamma EnrolRates_S * After_{it} + \boldsymbol{\beta}_1\mathbf{X}_{it} + \boldsymbol{\beta}_2\mathbf{W}_{ht} + \boldsymbol{\beta}_3\mathbf{C}_t + \boldsymbol{\beta}_4\mathbf{S} + \boldsymbol{\beta}_5\mathbf{S} * \mathbf{Y})$$

Results from this specification are presented in column 1 of Table 6.

of equation (3) for specifications analogous to those used in binary Probit. The treatment effects are significant at conventional levels of significance except the effect from LPM in specification A1 is not tightly estimated (p-value of 0.139).

Focusing on preferred specification in column 4, the Probit model indicates that children in the treated states are 21 percentage points less likely to enroll in grade 1 by age 7 relative to children in the control states during the pre-program period, and the effect is statistically significant at 1 percent level. This evidence supports the argument that there was pre-existing difference in the timing of enrollment in primary school between children in the treated and control states prior to the education program, where children in the treated states were less likely to enroll in primary school at the legal enrollment age relative to those in the control states.

Table 4: Probit DID estimates of the effect of education reform on on-time school enrollment (Observations 2368)

Dependent Variable: Binary Indicator for Enrollment in Grade 1 by Age 7

Variables	(1)	(2)	(3)	(4)	(5)
Treated*After	0.1887***	0.2146***	0.2197***	0.2133***	0.2305***
	(0.068)	(0.060)	(0.065)	(0.065)	(0.064)
State fixed effects	Yes	Yes	Yes	Yes	Yes
State-by-year fixed effects	Yes	Yes	Yes	Yes	Yes
Individual characteristics	No	Yes	Yes	Yes	Yes
Household characteristics	No	No	Yes	Yes	Yes
Locality & related Characteristics	No	No	No	Yes	Yes
Unemployment rate	No	No	No	No	Yes
Minus Log Likelihood	1273	1100	1012	1001	1001
AIC	2584	2250	2086	2075	2075

Statistical significance: *** p < 0.01, ** p<0.05, * p<0.10

Probit estimates of the treatment effect are based on standard errors clustered by enumeration area, the primary sampling units.

The control groups are individual-level characteristics (a binary indicator for gender, birth order, mother's and father's age and years of schooling), household-level characteristics (household size, binary indicators for whether a household has piped water, electricity, pit latrine, land, and farm animal), and locality-level and related characteristics (i.e., proportion

of households with piped water, electricity, pit latrine, land, farm animal, and urban dummy for location of residence), and state unemployment rate.
The post treatment year dummy (After) and the treatment variable (treated) are also included in each specification.
AIC denotes Akaike Information Criterion.

The average treatment effect of the interaction term is about 0.21 in Probit model and 0.35 in LPM (specification in column 4 of Table A3 in the Appendix). This suggests children in the treated state are 21 and 35 percentage points more likely to enroll on time relative to those who live in the control state after the program has occurred. Since the data set for the analysis is available only for two years, we are unable to formally test for common trend assumption. However, the specifications control for pre-existing differences in the timing of enrollment between children in the treated and control states; the time trend, i.e., the change in the timing of enrollment overtime due to other factors; observable individual, household, and community-level characteristics; state fixed effect; and state-by-year fixed effect. Hence, this effect is attributed to the education program, and it can be interpreted as the “true” average treatment effect.

Even if Table 4 documents positive and significant average treatment effect, it is crucial to examine the distribution of the treatment effect in non-linear models such as Probit since marginal effect is not constant in non-linear models. Figure 3, hence, presents the histogram and kernel density of the treatment effect based on the results reported in column 4 of Table 4 with average treatment effect of 0.21. The figure clearly shows that the estimated treatment effect is always non-negative and goes up well above 40 percentage points, suggesting large and positive treatment effect. The histogram and kernel density of the treatment effect is also plotted separately for rural and urban samples (see Figures B.1 and B.2 in Appendix B) to see if there is any difference in the treatment effect between rural and urban samples. The figures show strong and positive treatment effect both for urban and rural samples. However, the distribution for rural sample is right skewed while that for the urban sample is left skewed.

Figure 3: Historgram and Kernel Density Estimate of the Treatment Effect

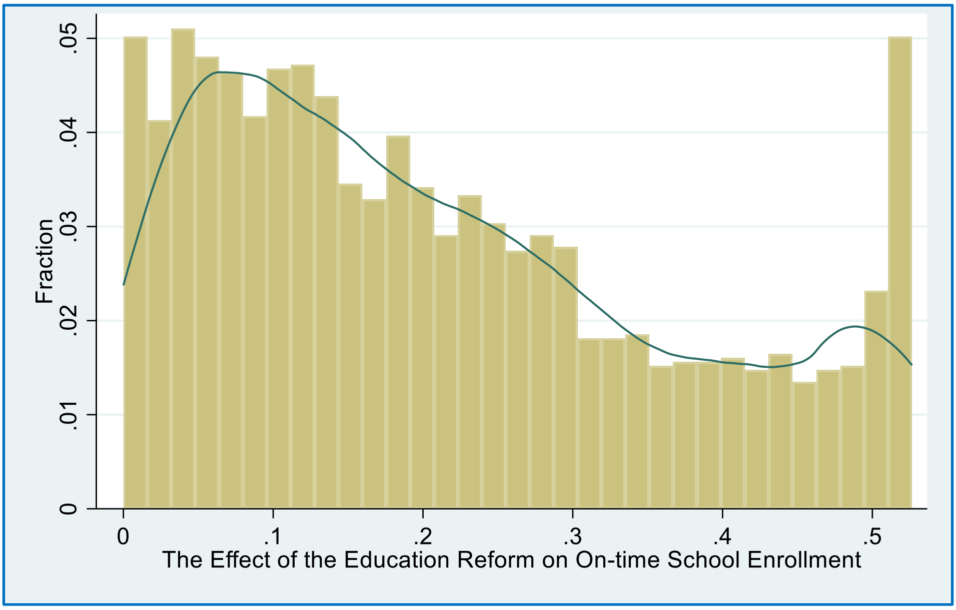

Alternative specifications and robustness check

The specifications presented in Tables 4 and A3 do not control for family income. This is because information on family income is not collected in the WMS data. Fortunately, however, detailed information on family income and expenditure is gathered in a supplementary survey called Household Income, Consumption, and Expenditure Survey (HICES), which is also administered by the Ethiopian Central Statistical Agency. HICES collects information on a subset of households that are surveyed in WMS, and it is usually conducted in the same year as the WMS. Using households sampled both in the HICES and WMS, we re-estimate equation (3) both by including and excluding household expenditure in the regression model.

Table 5 presents the results from different specifications using the restricted sample (i.e., households observed both in the HICES and WMS), and hence has relatively smaller sample size.

Table 5: Probit DID estimates of the effect of education reform on on-time school enrollment (Restricted Sample, Observation 1435)

Dependent Variable: Binary Indicator for Enrollment in Grade 1 by Age 7

Variables	**(1)**	**(2)**	**(3)**
Treated*After	0.2781***	0.2760***	0.2776***
	(0.090)	(0.090)	(0.090)
Log of HH expenditure	No	Yes	No
Log of community expenditure	No	No	Yes
State fixed effects	Yes	Yes	Yes
State-by-year fixed effects	Yes	Yes	Yes
Individual characteristics	Yes	Yes	Yes
Household characteristics	Yes	Yes	Yes
Locality & related Characteristics	Yes	Yes	Yes
Minus Log Likelihood	632	632	632
AIC	1337	1338	1339

Statistical significance: *** $p < 0.01$, ** $p<0.05$, * $p<0.10$

Probit estimates of the treatment effect are based on standard errors clustered by enumeration area, the primary sampling units.

The control groups are individual-level characteristics (a binary indicator for gender, birth order, mother's and father's age and years of schooling), household-level characteristics (household size, binary indicators for whether a household has piped water, electricity, pit latrine, land, and farm animal), and locality-level and related characteristics (i.e., proportion of households with piped water, electricity, pit latrine, land, farm animal, and urban dummy for location of residence).

The post treatment year dummy (After) and the treatment variable (treated) are also included in each specification.

To make comparison of results from different specifications (that control for household expenditure and that do not) straight forward, column 4 of Table 4 is re-estimated for the restricted sample, and the results are presented in column 1 of Table 5. Column 2 of Table 5 presents the results from a specification that controls for household expenditure. Controlling for household expenditure changes neither the magnitude nor the significance of the average treatment effect of the interaction term.

The coefficient estimate (not reported here) of the household expenditure itself, on the other hand, is positive, but not significant. It is insignificant may be because the specification controls for parental years of schooling and household assets and amenities, which are generally good controls for families' socioeconomic status.

If higher income families self-select themselves to live at closer proximity to schools and they are more likely to enroll their children in primary school on time regardless of their proximity to school, then household income or expenditure is endogenous and bias the results. The program was explicitly designed to make primary schools more accessible to households in rural areas and underserved localities. In this setting, bias from this type of selection is less likely since the program exogenously allocates new schools across households. If higher income families somehow managed to influence policy makers to build more schools in their locality or higher income families move to areas that received more school construction, then household expenditure is endogenous and biases the results. To mitigate potential endogeneity of household expenditure, we aggregated household expenditure at enumeration area a primary sampling unit which is typically equivalent to a village or urban neighborhood level and estimated equation (3). The results are depicted in column 3 of Table 5. The average treatment effect under this specification is again similar to those presented in columns 1 and 2 both in magnitude and significance. The similarity of the results reinforces the argument that relatively rich communities were less likely to influence policy makers to build more schools in their communities. It also suggests there is no evidence that high income families moved to areas that received more school allocation.

One of the identifying assumptions in the difference-in-differences model is the economic growth rate in the treated and control states do not vary systematically over time. In reality, however, states in the two groups may experience different growth rates. Thus, the estimates could potentially confound the effect of the program with the effect of the differential growth rate on the timing of enrollment that would have been observed even in the absence of the program. Thus, we present a specification that controls for state level unemployment rate in column 5 of Table 4. Information on unemployment rate is obtained from the 1994 and 2007 Ethiopian census. In this specification, the average treatment effect has increased by about 4 percentage points relative to the basic specification.

If we expect states with relatively higher growth rate (or lower unemployment rate) make schools relatively more accessible to their residents in the absence of the program, and if we assume the program targets building more schools in states with lower growth rate, then comparison of the average treatment effect in the basic specification and the one that controls for differences in economic growth rate implies that the program help children who live in lower-growth-rate states to catch up with those in high-growth-rate states in terms of enrolling in primary school on time.

Table 6 gives results from specifications that involve alternative outcomes and treatment status as appropriate. Column 1 of Table 6 presents results from a model that replaces a binary indicator (for treated states) by a continuous measure of pre-program state level primary school enrollment rate. One advantage of using a continuous primary school enrollment rate variable, rather than a binary indicator, is it makes use of all the available information and hence the treatment effect is more precisely estimated. Besides, it is more robust to the risk of arbitrarily grouping states into treatment and control groups. Prior studies employ a similar strategy to estimate treatment effect. For instance, Miller (2012) used pre-reform insurance rate to investigate the effect of the 2006 Massachusetts health reform on emergency room visits. In this continuous treatment specification, Treateds in equation (3) is replaced by pre-program state level primary school enrollment rate, say EnrolRates.

The dependent variable for specifications 1 through 3 is a binary indicator for enrollment in grade 1 by age 7. The response variable for Spec 4 is the logarithm of age at enrollment in grade 1. Estimates were obtained using probit (specs 1-3) and OLS (Spec 4).

The control groups are individual-level characteristics (a binary indicator for gender, birth order, mother's and father's age and years of schooling), household-level characteristics (household size, binary indicators for whether a household has piped water, electricity, pit latrine, land, and farm animal), and locality-level and related characteristics (proportion of households with piped water, electricity, pit latrine, land, farm animal, and urban dummy for location of residence).

Table 6: DID estimates of the effect of education reform on school enrollment: continuous and binary treatments (N = 2368)

Variable	Dependent variable: Enrollment Status			ln(AgEnrollment)
	(1)	**(2)**	**(3)**	**(4)**
Enrollment rate*After	-0.0246**			
	(0.012)			
School distance*After		0.0383		
		(0.048)		
Planned # schools*After			0.0005	
			(0.0003)	
Treated*After				-0.1973***
				(0.069)
State fixed effects	Yes	Yes	Yes	Yes
State-by-year fixed effects	Yes	Yes	Yes	Yes
Individual Characteristics	Yes	Yes	Yes	Yes
Household characteristics	Yes	Yes	Yes	Yes
Locality & related covariates	Yes	Yes	Yes	Yes
R-squared				0.382
Minus log-likelihood	1001	1001	1001	
AIC	2075	2075	2075	

Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$
Standard errors in parentheses are clustered by enumeration area, the primary sampling unit.
The dependent variable for specifications 1 through 3 is a binary indicator for enrollment in grade 1 by age 7. The response variable for Spec 4 is the logarithm of age at enrollment in grade 1. Estimates were obtained using probit (specs 1-3) and OLS (Spec 4).
The control groups are individual-level characteristics (a binary indicator for gender, birth order, mother's and father's age and years of schooling), household-level characteristics (household size, binary indicators for whether a household has piped water, electricity, pit latrine, land, and farm animal), and locality-level and related characteristics (proportion of households with piped water, electricity, pit latrine, land, farm animal, and urban dummy for location of residence).
The post treatment year dummy (After) and the relevant treatment variable are also included in each specification.

In this model, the estimate of the average marginal effect of *EnrolRates* can be interpreted as the change in the probability of enrollment in grade 1 by age 7 for a one percent change in the pre-program enrollment rate. We find that children who lived in states with one percent higher pre-program primary school enrollment rate were about 1.9 percentage points (not reported in Table 6) more likely to enroll in school on time, reaffirming the pre-existing difference on the timing of enrollment across children that live in states with different pre-program primary school enrollment rate. On the other hand, the average treatment effect is estimated to be -0.025. This suggests that, on average, children that lived in a state with one percent higher pre-program enrollment rate were 2.5 percentage points less likely to enroll in primary school on time. Thus, the program has caused children that live in states with lower pre- program enrollment rate to enroll in school on time.

One of the potential concerns of our identification strategy is that the variable we chose (i.e., pre-program primary school enrollment rate) to measure the variation in the intensity of the impact of the program across districts could be arbitrary. As discussed above, we chose pre-program primary school enrollment rate because the program's major objective was to increase primary school enrollment rate from its pre-program level of 30 percent to at least 50 percent. However, we want to check if our results are sensitive to the choice of alternative variables. Columns 2 and 3 of Table 6 present the results from this exercise.

In column 2 of Table 6, we present results from a specification where we use continuous measure of state-level pre-program distance to primary school (instead of pre-program primary school enrollment rate with its result reported in column 1 of Table 6). Since the program intends to make schools more accessible to previously neglected areas, we should expect children who lived in states with longer average distance to primary school before the program to more likely enroll on time relative to those that lived in states with shorter distance to primary school. The magnitude of the average treatment effect of 0.038 can be interpreted as children that lived in a state with an average of 1 km longer pre-program distance to primary school were about 3.8 percentage points more likely to enroll in primary school on time as a result of the intervention/program. However, the estimate is not statistically significant.

Alternatively, we used the number of schools that the federal government has planned to build in each state as part of the program.27 Again, the government has planned to build more schools in underserved areas and hence we should expect the impact of the program to be stronger in areas where the government has planned to build more schools. Under this specification (which is reported in column 3 of Table 6), the average treatment effect is estimated to be 0.0005. This treatment effect

[27] Ideally, we would want to use the variation in the actual number of schools built across districts as a result of the program. Since we do not observe the actual number of schools built in each state, we resort to using the number of schools the government has planned to build in each state.

suggests that building additional 100 primary schools have, on average, increased the probability of on-time enrollment by about 5 percentage points. This effect is not tightly estimated.

Overall, the results reported in Table 6 show that our results are largely not sensitive to the choice of different measures of the variation in the intensity of the impact of the program across districts in Ethiopia.

By how much has the program decreased age at enrollment?

The results presented above support the argument that the program has increased the probability of enrollment in grade 1 at the legal enrollment age. It is, therefore, interesting to investigate by how much the program has decreased age at enrollment. Column 4 of Table 6 presents the estimates of this experiment where the dependent variable is the natural logarithm of age at enrollment. Again, results from the basic specification of the probability model are reported in column 4 of Table 4 with the estimated average treatment effect of 0.213. As discussed earlier, this suggests that children in the treated state are 21.3 percentage points more likely to enroll in primary school on time. Column 4 of Table 6, on the other hand, depicts the results of the log duration regression, where the dependent variable is the natural logarithm of age at enrollment. The results reported in column 4 of Table 6 show that the average treatment effect is -0.197, implying that the difference in age at enrollment between children in the treated and control states has decreased by 19.7 percent as a result of the education program. Remember that children in the treated states, on average, enroll in primary school 1.57 years later than those in the control states before the education program (see Table 3). Hence, the program has decreased age at enrollment in grade 1 by 0.31 ($1.57 * 0.197$) years or 3.72 months.

6. Conclusion

In recent years, many governments in developing countries have attempted to achieve universal primary education through a largescale construction of primary schools. The majority of the studies on primary education in developing countries focus on enrollment rates, without considering the timing of enrollment. Delaying primary school enrollment beyond the legal enrollment age, however, is more of a norm than an exception in these countries. Prior studies have documented that delaying enrollment is costly as it, for instance, decreases an individual's lifetime wealth (Glewwe and Jacoby, 1995), and it increases both school dropout and grade repetition rates (Wils, 2004). Though delayed enrollment is widely observed in developing countries and there is a high cost associated with it, the literature on the topic is limited, and we have a limited understanding of why children delay enrollment in primary school. This paper attempts to fill the gap in the literature by investigating the effect of access to primary school on the timing of enrollment in primary school.

Identifying the causal effect of access to primary school on the timing of enrollment is complicated by endogeneity of access to primary school. For instance, parents who choose to live at closer proximity to school may have strong taste for education and enroll their children in school on time regardless of proximity to school. To mitigate biases due to endogeneity of access to school, we exploit the education reform in Ethiopia as exogenous source of variation in access to primary school. Then, we estimated difference-in-differences model where the dependent variable is a binary indicator for enrollment in primary school by age 7, the legal enrollment age in Ethiopia, and the natural logarithm of age at enrollment in primary school.

The average treatment effect is estimated to be between 0.21 and 0.35, suggesting the probability the child enrolls in primary school on time has increased by between 21 and 35 percentage points as a result of the education reform. The log duration DID regression (where the dependent variable is the natural logarithm of age at enrollment), on the other hand, suggests that the reform has decreased age at enrollment in grade 1 by about 4 months. These estimates highlight the important role access to school plays in inducing parents to enroll their kids in primary school at the legal enrollment age.

The findings reported here are important as they show that, in Ethiopia, education intervention has been effective in decreasing age at enrollment in primary school. The intervention was meant to increase primary school enrollment, but it also induced households to enroll their children in primary school at a relatively younger age. The Ethiopian government provides free primary education. Thus, households do not have to pay for tuition. Households, however, still have to incur other costs related to school attendance, including the child's opportunity cost of time in terms of forgone family income from child labor.

Making schools accessible to poor households would decrease the time the child spends walking to school, and hence decreases the opportunity cost of school attendance. Moreover, accessibility induces physically weaker children to attend school since it decreases the physical strength needed to walk the distance to school. Policy makers, thus, should also consider improving communication networks and public transport as alternative/additional ways to encourage households to enroll their kids in primary school on time. At this point, it is worth mentioning that there is inequality in access and on-time enrollment in Ethiopia where girls, kids with disabilities, and kids from rural areas and emerging states are at disadvantage. Thus, policy choices that focus on narrowing down inequalities across groups should be given a priority.

Providing quality education has increasingly become a challenge in Ethiopia given the scale of expansion and increased access in the last couple of decades. However, the data still show that kids in Ethiopia do not enroll on-time at their legal school starting age. This has recently been exacerbated by COVID-19 and the on-going conflict throughout the country where thousands of schools have been destroyed and millions of kids are forced to be out of school at the moment, see, for example, Ministry of Education (2022). A country like Ethiopia should still need to work on both improving access and quality; of course, along with improving equity and finding sustainable ways of financing the education sector. We, thus, believe this study will inform not only Ethiopian policymakers but also other policymakers in similarly resource-constrained countries.

References

Barro, Robert J. and Jong-Wha Lee. (2001). "International data on educational attainment: updates and implications," Oxford Economic Papers, July 2001, 53 (3), 541–563.

Bommier, Antoine and Sylvie Lambert. (2000). "Education Demand and Age at School Enrollment in Tanzania," *The Journal of Human Resources*, January 2000, 35 (1), 177–203.

Chen, Qihui. (2015). "Ready for School? Impacts of delayed primary school enrollment on children's educational outcomes in rural China," *International Journal of Educational Development*, 2015, 45, 112-128.

De Vreyer, Philippe, Sylvie Lambert, and Thierry Magnac. (1999). "Educating Children: A Look at Household Behavior in Cote d' Ivoire," EPEE #99-13, 1999.

Glewwe, Paul and Hanan G. Jacoby. (1995). "An Economic Analysis of Delayed Primary School Enrollment in a Low-Income Country: The Role of Early Childhood Nutrition," *The Review of Economics and Statistics*, February 1995, 77 (1), 156–169.

Jacoby, Hanan G. (1994). "Borrowing Constraints and Progress Through School: Evidence from Peru," *The Review of Economics and Statistics*, February 1994, 76 (1), 151–160.

Lloyd, Cynthia B. and Ann K. Blanc. (1996). "Children's Schooling in sub-Saharan Africa: The Role of Fathers, Mothers, and Others," Population and Development Review, June 1996, 22 (2), 265–298.

McEwan, Patrick J. (2013). "The impact of Chile's school feeding program on education outcomes," Economics of Education Review, February 2013, 32, 122–139.

Miller, Sarah. (2006). "The effect of insurance on emergency room visits: An analysis of the 2006 Massachusetts health reform," *Journal of Public Economics*, December 2012, 96 (11-12), 893–908.

Ministry of Education. (2022). "National Study on the Magnitude of Out-of-School Children in Ethiopia," Ministry of Education.

Moyi, Peter. (2010). "Household characteristics and delayed school enrollment in Malawi," *International Journal of Educational Development*, May 2010, 30 (3), 236–242.

Partnership. (1999). "Short stature and the age of enrolment in Primary School: studies in two African countries. The Partnership for Child Development," Social science & medicine (1982), March 1999, 48 (5), 675–682.

Puhani, Patrick A. (2012). "The treatment effect, the cross difference, and the interaction term in nonlinear 'difference-in-differences' models," *Economics Letters*, 2012, 115(1), 85-87.

Seshi-Nasser, Hellen A. and Abena D. Oduro. (2016). "Delayed primary school enrollment among boys and girls in Ghana," *International Journal of Educational Development*, 2016, 49, 107-114.

Todd, Jessica E. and Paul Winters. (2011). "The Effect of Early Interventions in Health and Nutrition on On-Time School Enrollment: Evidence from the Oportunidades Program in Rural Mexico," Economic Development and Cultural Change, April 2011, 59 (3), 549–581.

Wils, Annababette. (2004). "Late entrants leave school earlier: evidence from Mozambique," International Review of Education, 2004, 50 (1), 1737.

World Bank. (1998). "Program Appraisal Document on a Proposed International Development Association Credit in The Amount of US $100 Million Equivalent to the Federal Democratic Republic of Ethiopia for the Education Sector Development Program," World Bank, 1998.

World Bank. (2005). "Implementation Report on a Loan/Credit/Grant in the Amount of SDR 74,300,000 (US $ 100 Million Equivalent) to the Federal Democratic Republic of Ethiopia for an Education Sector Development Project," World Bank, 2005.

Appendix

Appendix A. Additional Tables

Table A.1: List and Description of Variables Used in Estimation

Variable	Description
Dependent variables	
Enrolled in grade 1 by age 7 (yes=1)	=1 if a child is enrolled in grade 1 at age ≤ 7; 0 otherwise
Age at enrollment	Age in years by the time a child enrolled in grade 1
Independent variables	
Treated	=1 if state level primary school enrollment rate is ≤ 30; 0 otherwise
After	=1 if year=2004; 0 otherwise
Girl	=1 if a child is a girl; 0 otherwise
Birth order	=1 for a first-born child, =2 for a second-born child, etc
Household size	Total number of people who live in the household
Dad's years of schooling	Highest grade completed by the father
Mom's years of schooling	Highest grade completed by the mother
Dad's age	Father's age in years
Mom's age	Mother's age in years
HH has piped water	=1 if the household has piped water; 0 otherwise
HH has electricity	=1 if the household has electricity; 0 otherwise
HH has pit latrine	=1 if the household has pit latrine; 0 otherwise
HH owns land	=1 if any member of the household owns any land holdings regardless of how the land is used; 0 otherwise
HH owns farm animal	=1 if the household owns farm animals 0 otherwise
Proportion of hhs with piped water	Proportion of households with piped water
Proportion of hhs with electricity	Proportion of households with electricity
Proportion of hhs with pit latrine	Proportion of households with pit latrine
Proportion of hhs with land	Proportion of households that owns land
Proportion of hhs with farm animal	Proportion of households that owns farm animal
Log (Household expenditure)	Log of annual total household expenditure in 2005 prices
Unemployment rate	State level unemployment rate
Urban area (yes=1)	=1 if the household is located in urban area; 0 otherwise

Table A.2: Fraction of Children Enrolled in Grade 1 by Age 7 by Treatment Group and Time (Observations 2368)

Group	Before the Change	After the Change	Time Difference
Treated	0.228***	0.224***	-0.004
	(0.020)	(0.010)	(0.022)
Untreated	0.476***	0.365***	0.112*
	(0.046)	(0.041)	(0.062)
Group Difference	-0.248***	-0.141***	0.107
	(0.050)	(0.042)	(0.081)

Notes: Standard errors are given within parentheses. The standard error associated with the treatment effect (highlighted) is clustered by enumeration area, the primary sampling unit.
Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Table A3: LPM DID Estimates of the Effect of Education Reform on On-time School Enrollment (Observations 2368)

Dependent Variable: Binary Indicator for Enrollment in Grade 1 by Age 7

Variables	(1)	(2)	(3)	(4)	(5)
Treated*After	0.3449	0.3519**	0.3469**	0.3514**	0.3253*
	(0.233)	(0.165)	(0.176)	(0.173)	(0.169)
State fixed effects	Yes	Yes	Yes	Yes	Yes
State-by-year fixed effects	Yes	Yes	Yes	Yes	Yes
Individual characteristics	No	Yes	Yes	Yes	Yes
Household characteristics	No	No	Yes	Yes	Yes
Locality & related Characteristics	No	No	No	Yes	Yes
Unemployment rate	No	No	No	No	Yes
R-squared	0.030	0.171	0.230	0.239	0.239

Statistical significance: *** $p < 0.01$, ** $p<0.05$, * $p<0.10$
LPM/OLS estimates of the treatment effect are based on standard errors clustered by enumeration area, the primary sampling unit.

The control groups are individual-level characteristics (a binary indicator for gender, birth order, mother's and father's age and years of schooling), household-level characteristics (household size, binary indicators for whether a household has piped water, electricity, pit latrine, land, and farm animal), and locality-level and related characteristics (i.e., proportion of households with piped water, electricity, pit latrine, land, farm animal, and urban dummy for location of residence), and state unemployment rate.

Appendix B. Additional Figures

Figure B.1: Histogram and Kernel Density Estimate of the Treatment Effect, Rural Sample

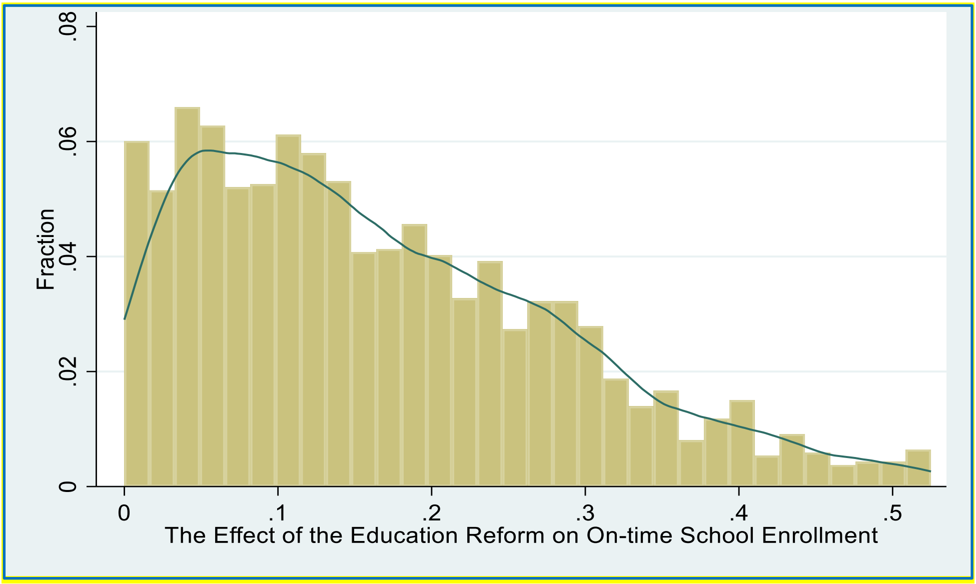

Figure B.2: Histogram and Kernel Density Estimate of the Treatment Effect, Urban Sample