Fertility Responses to Schooling Costs: Evidence from Uganda's Universal Primary Education Policy *

Alfredo Burlando† *and Edward Bbaale*‡

July 21, 2020

Abstract

There is some evidence that access to schooling reduces fertility along the intensive margin in developing countries, but the transmission channels are not well understood; most education interventions impact financial costs, access, and school quality. We isolate the specific effect of child school fees on maternal fertility by studying a 1997 schooling reform in Uganda, in which the government abolished elementary school fees for up to four children per household. Families with more school-aged children were required to pay fees for the additional ones. We demonstrate that eligibility limits reduced births: women with more than four children in 1997 were 4.2 percentage points less likely to give birth to an additional child in the subsequent two years. We provide evidence that this result is driven by UPE eligibility limits, not by other factors associated with the policy or other secular changes. Fertility effects are persistent over time and survive the lifting of the eligibility limit in 2003. The policy is also associated with a temporary reduction in the desired fertility, with mothers preferring having four children or fewer during the policy period than before.

*We are grateful for helpful comments from Carly Urban, Andrew Foster, Rachel Heath, Caroline Weber, Jamie McCasland, Munir Squires, Fernando Aragon, Todd Pugatch, and participants at the Northwest Development Workshop in 2017 and the Population Association of America Annual Meeting in 2018, as well as from seminar participants at the University of Oregon and the University of Washington.

†Department of Economics, University of Oregon, Eugene, OR 97403-1285, USA, burlando@uoregon.edu.

‡School of Economics, Makerere University, Plot 51 Pool Road, Kampala, Uganda, eddybbaale@gmail.com.

JEL classification: J13, I25, O12

Keywords: Fertility, Education, Quantity and Quality of Children, Universal Primary Education, Uganda, Africa.

1 Introduction

The literature documenting a negative relationship between schooling and fertility in developing countries is very long and robust. Many studies correlate increases in overall schooling to reductions in total fertility rate (e.g., Ainsworth et al., 1996; Lloyd et al., 2000; Kravdal, 2002), while others use policies that improve access to schooling as shocks to fertility (Osili and Long, 2008; Behrman, 2015; Keats, 2018; Chicoine, 2020). Despite the existence of this empirical relationship, the mechanisms through which increases in schooling lead to fertility reductions are not completely understood. Education can delay the time of first marriage, shape preferences regarding household size and household composition, and change the opportunity cost of a woman's time. At the same time, an increase in educational opportunities for children can change the incentives of parents to invest in their children, and lead to a reallocation of household resources toward having fewer, more educated children as hypothesized by Becker (1960) and Becker and Lewis (1973). Separating these possible pathways is important but not always possible. For instance, the literature that links exogenous shifts to educational opportunities caused by universal primary education (UPE) reforms cannot separate the roles of reducing monetary costs to schooling, improvements in access, changes in desired fertility levels, or changes in the opportunity cost of fertility for mothers. Understanding the contribution of each pathway to the overall effect matters because different schooling policies may not affect all channels equally.

This paper aims to shed light on one specific mechanism: the role of monetary schooling costs on the fertility decisions of parents who incur those costs. Understanding this mechanism is important since education policies in low- and middle-income countries often aspire to reduce these costs. Many countries have experimented with: conditional cash transfer programs (CCTs), which are partially used to defray schooling costs; voucher programs, designed to improve access to private schools; school feeding programs and school uniform subsidies, which reduce ancillary household costs; and the elimination of school fees for public education. While each of these programs has unique features and requires careful consideration, all relax household budget constraints in one way or another.[1] In addition, this may be a particularly important channel in sub-Saharan Africa, where fertility rates are high and demand for schooling is quite elastic.

In this paper, we more clearly isolate the effects of schooling costs in a UPE reform with some unusual features. In 1997, the Ugandan government implemented a schooling reform that abolished elementary school fees but limited eligibility for free schooling to four children per household. Households were required to pay fees for each additional child attending primary school regardless of when they attended (i.e., years after the older siblings) or where they enrolled. In 2003, this controversial system was abolished and all children gained access to free primary school. Conceptually, the 1997 reform created a system of prices for schooling that differed by household composition. In households with fewer than four school-aged children ("small" households), all children faced the same (lower) schooling costs. Among households with more than four children, the cost of schooling for the marginal child remained as high as the prereform period, while the (inframarginal) cost of the first four children declined. Thus, this policy generated different schooling costs for marginal children who attended the same school and lived in the same community but belonged to households that differed in composition.

We study the impact of variation in costs of schooling on the fertility of mothers by exploiting this relationship between costs and household structure. In a generalized difference-

[1] Fertility impacts of CCT programs remain an area of active research (Todd and Wolpin, 2006; Stecklov et al., 2007; Baird et al., 2011; Fiszbein and Schady, 2009)

in-differences framework that allows us to control for community-wide shocks to schooling and fertility, we compare changes in births for mothers who, at the onset of the policy, had four children or more of elementary school age, relative to mothers who had fewer children. We find that, following the policy, these mothers reduced the likelihood of another birth within 46 months from the onset of the policy by 4.2 percentage points; this is a meaningful drop, as it equals 8 percent of the average likelihood of an additional birth. We then run a number of robustness checks to ensure our results are driven by the UPE policy. We first demonstrate that fertility responses to the policy are sharply different between mothers with three children and mothers with four children, consistent with the latter but not the former being constrained by the policy. We also run placebo regressions that vary both the timing of the policy and the placement of the eligibility cutoff, and find that fertility responses appear only around the time of the policy and at the expected cutoff. We are thus able to rule out obvious violations of the parallel trend assumption, including the presence of long-term downward trends in fertility. We conclude our analysis by studying the elimination of the eligibility rule in 2003 and showing that reductions in fertility are persistent.

Our results are consistent with a quality response associated with the quantity-quality tradeoff model. The results are also consistent with an alternative explanation: a change in social norms around acceptable or ideal family sizes, due to the policy implicitly favoring families with four or fewer children. While social norms around fertility tend to change slowly over time (Munshi and Myaux, 2006), and the policy was justified on budgetary concerns only, we do provide some intriguing evidence that the ideal number of children was lower in 2000 than either before or after the policy, a fact that can support this explanation. Our results are not driven by more complex general equilibrium mechanisms, as our within-community framework controls for local equilibrium effects of the policy such as the learning environment.

Our paper contributes to our understanding of the importance of schooling policy on fertility decisions in developing countries. A strand of the existing literature studies UPE reforms (Osili and Long, 2008), including Uganda's (Behrman, 2015; Keats, 2018). While the literature finds that UPE reduces desired and actual fertility, this is achieved through a number of pathways. Others have looked at other programs such as CCTs, which provide income to families who send children to school. By linking transfers to the household to the number (and school attendance) of children, these programs have a priori an ambiguous effect on fertility. The literature has found no effects of CCT programs on fertility in Mexico (Todd and Wolpin, 2006; Stecklov et al., 2007), although more generous programs seem to induce households into higher fertility (Stecklov et al., 2007). To our knowledge, there is no comparable set of estimates from sub-Saharan Africa as the scale of CCT programs there is too small to address the question.[2]

This paper also contributes more broadly to a long quantity-quality literature that began with Becker and Lewis (1973). There is little evidence for quantity-quality trade-offs in developed countries (Doepke, 2015), but the evidence from developing countries is much stronger (Schultz, 1997). The demographic literature finds a strong correlation between fertility and schooling (Ainsworth et al., 1996; Kravdal, 2002; Lloyd et al., 2000) but cannot disentangle the feedback effects between schooling and fertility.[3] Our paper contributes specifically by highlighting a previously unexplored feature of this model. We show that reductions in schooling costs, which yield ambiguous effects on fertility, can be decomposed into reductions in marginal and average costs; that reductions in average cost conditional on high marginal cost reduce fertility. The Uganda policy allow us to isolate this fertility response, and our results indicate that fertility responds strongly to quality.

[2] An exception is Baird et al. (2019), who find significant fertility responses to CCT among direct beneficiaries in Malawi. As in the UPE fertility literature, a number of channels can explain these fertility impacts.

[3] Another important strand of the quantity-quality literature studies how shifts in fertility affect schooling. Much of that evidence comes from exogenous changes to quantity caused by fertility policies or by twinning (Rosenzweig and Zhang, 2009; Rosenzweig and Wolpin, 1980; Qian, Qian). While the policy studied here is a schooling policy, household size-based eligibility limits share some similarities with fertility policies.

The rest of the paper is organized as follows: Section 2 describes in greater detail the UPE policy implemented in Uganda in 1997 and adapts the quantity-quality framework to the setting. Section 3 explains the identification strategy and provides a description of the data used. Section 4 reports the main results of the paper. Section 5 concludes.

2 Background Information

2.1 Fertility trends in Uganda

Uganda is a high-fertility country with current total fertility rates (TFR) estimated around 6.2 live children per woman, above the sub-Saharan African average of 5.2 live children per woman (Kabagenyi et al., 2015). Indeed, Uganda has the tenth highest TFR in the world (Ariho et al., 2018). Historically, Uganda's fertility rates have fallen over time: they were between 7 live children per woman and 8 live children per woman in the eighties. An analysis of fertility trends by Kabagenyi et al. (2015) shows a slow, steady reduction in fertility between 1980 and 2011. There is no evidence of particular shifts in secular trends in the mid- to late nineties or in the early 2000s.

2.2 UPE reform in Uganda

The Ugandan government announced the UPE policy during the then president Museveni's reelection campaign speech in which he proposed to eliminate school fees, PTA fees, and building fees starting the following academic year, in January 1997. (Museveni, 1996). The objective of the policy was to "enable Ugandan children of school-going age (6-12) to enter and remain in school and complete the Primary Cycle of Education. This should be achieved as soon as possible but not later than the year 2003" (MoE&S, 1998). The government replaced the payment of school fees with a centrally allocated capitation grant–equivalent to UGX 5,000 per pupil for grades one through three, and UGX 8,100 per pupil in grades four to seven–provided to the school in monthly installments.

Crucially, both the presidential speech and subsequent policy explicitly and repeatedly limited the capitation grant (and the offer of free education) to four children per household. The policy applied to monogamous, polygamous, and single-parent families, and the four-child limit was defined as "once in a lifetime" (MoE&S, 1998). Thus, additional children were required to pay school fees.[4] Finally, despite the objective of increasing schooling for children aged 6 to 12, no specific limit on the age of the beneficiary child was imposed. As was typical in Ugandan schools, where students as old as 16 remained enrolled in primary school, the policy benefited children outside the targeted age cohort.[5]

It is important to highlight the fact that, as far as what we know from official documentation and experience, the policy was not presented to the public as an attempt to discourage large families. Rather, it appears that Museveni's intention at the time was to manage the trade-off between public expenditure and access for all households. His campaign speech said: "If you have more than four children, you will pay only their school fees but not PTA or building fees. In this way we shall be able to send as many children as possible to school" (Museveni, 1996). Given the large sizes of Ugandan households, the policy intended to manage schooling access and avoid huge enrollments. Nonetheless, enrollments increased tremendously and the limit of free schooling to four children proved controversial. In 2003, the government abolished eligibility limits.

The UPE policy included a number of other relevant interventions, namely school construction, school refurbishing, teacher housing construction, and teacher training and hiring. To finance the reform, the share of the education budget going to primary education increased from 40 percent in 1996 to 65 percent in 2004, and the overall education budget increased from 1.6 percent to 3.8 percent of GDP (Deininger, 2003). Finally, while there is

[4] An exception was given to orphans, who were entitled to free education.

[5] According to the DHS, in 1995, two thirds of children aged 16 were still in school. Of these, close to 25% were enrolled in primary school.

evidence that the implementation of the policy suffered from many problems, an analysis by Deininger (2003) found a significant reduction in school fees paid by households and an increase in the number of children attending school. He found that overall education spending per child enrolled in primary school (including private institutes) fell by almost 9,000 UGX, from 26,000 UGX pre-policy (reduction of 33%), and the proportion of children paying PTA fees fell from 96% pre-policy to 23% in 1999. The effect on the household limit was not analyzed. However, qualitative evidence also suggests the policy impacted class sizes, teachers' qualifications, and other elements related to the quality of schooling.

2.3 Conceptual framework

A natural way to understand how the UPE reform influenced parental fertility choices is to turn to a formal model of fertility and schooling such as the Becker and Lewis (1973) quantity-quality trade-off model, as its predictions (of a negative relationship between human capital and fertility) fit high fertility countries quite well (Doepke, 2015).

In that model, which we formally present in the appendix, households choose the number and quality (i.e., level of schooling) of their children. At the margin, the decision to have one more child depends on the shadow price of child quantity; this shadow price increases with the cost of raising the child and with the average level of schooling of all children in the household. Given this setup, changes to schooling costs (say, from a standard UPE policy that eliminates all student fees from all students) have ambiguous effects on overall fertility. On the one hand, a reduction in the cost of schooling reduces the cost of educating the marginal child, which encourages additional fertility; on the other hand, it encourages the household to increase the amount of schooling for all children, which discourages additional fertility. In addition, the income effect from lower schooling costs increase the demand for children (as long as they are a normal good).

The above prediction applies to a standard Universal Primary Education reform. In

Uganda, the reduction in costs applies only for the first four children in the household. This reduction is inframarginal for some households: the cost of educating the first four children is lowered, while for higher order children, including the marginal child, it is unchanged. Thus, the shadow price of the marginal child unambiguously increases, because child quality is increased. The Ugandan policy should therefore induce a strong quality response.[6]

Naturally, the UPE policy might affect fertility through other mechanisms. For example, the policy could have changed social norms around acceptable or ideal family sizes, by implicitly endorsing families with four children or fewer. Indeed, existing evidence does indicate that changes in information about social norms can lead to updated beliefs about those norms, which in turn affect household decisions (Bursztyn et al., 2018). However, social norms around fertility seem to change slowly (Munshi and Myaux, 2006).

3 Empirical strategy

Data Our data consist of completed mother-level birth histories from the Demographic Health Surveys (DHS) from 1995 and 2000-2001. The 1995 wave will provide fertility and household information for the period preceding the UPE policy. The 2000-2001 wave will provide information on fertility and households shortly after the policy was implemented.[7] We will also make use of the 2006 survey, which took place almost a decade after implementation and three years after limits on tuition waivers were lifted.

Empirical strategy Our baseline empirical strategy is based on variation due to the timing of the policy implementation occurring after 1997 and the composition of the household. Our

[6] The original quantity-quality model assumes a strong equitability motive where all children receive the same amount of education. This assumption can be relaxed: as long as parents have some taste for equality, lower schooling costs for the first four children increases quality for the marginal child, making her more costly. A positive impact on fertility could exist if parents substitute away from the last child (Becker and Tomes, 1976). Quality responses can be elicited through other mechanisms: for example, if parents of large families are disproportionally more likely to move children to private schools.

[7] Specifically, the collection of data in 2000 began in October and continued through March 2001.

unit of analysis is a woman m aged 15-49 residing in community c at the time of the 1995 or 2000 DHS survey. Using each woman's birth records at the time of interview in 2000 or 2001, we construct a measure of the number of children already alive in December 1996 (right before the UPE policy was started). We then generate a variable $LargeHousehold$ that indicates whether the woman had at least four children aged 0-16 in December 1996 and was therefore potentially subject to the limit on additional births. Finally, we construct the dependent variable $AdditionalBirth$ as an indicator equal to one if at least one more child was born between January 1997 and October 2000 (corresponding to the start of the DHS). The fertility period thus covers 46 months.[8]

Having defined outcome and household composition for women in the treatment period, we replicate the same strategy for women observed in the control period, that is, who were interviewed in 1995. We calculate the number of live offspring in December 1990, identify mothers with more than four children at that time with "large household" indicator, and identify mothers who gave birth again between January 1991 and October 1994 by the indicator variable $AdditionalBirth$.[9] A schematic of the approach is shown in Appendix Figure A1.

Our estimation strategy relies on the following generalized difference-in-differences linear probability model:

$$AdditionalBirth_{mtc} = \alpha PostUPE_t \times LargeHousehold_m + \sum_{j=1}^{10} \gamma^j Sib_m^j + X_m\beta + \delta_{tc} + \epsilon_{mtc}. \quad (1)$$

The difference-in-differences estimate α identifies the level shift caused by the interaction

[8] Note that our definition of the household, for the purposes of this paper, includes only children born to a mother. Household structures as reported in the DHS are more complex and flexible, however, they can be manipulated as a response to the policy.

[9] Note that DHS interviews started in March 1995 and October 2000. Thus, there is a slight discrepancy between the recall periods for the pre- and post-UPE groups (i.e., we require a slightly longer recall period for the pre-UPE group). We prefer evaluating fertility over the period starting in January in both pre- and post-UPE periods so as to avoid seasonality differences. Different evaluation periods in the pre-UPE period do not change our results.

between a dummy variable identifying women interviewed after the reform ($PostUPE_t$) and a dummy variable identifying households with four or more children aged between 0 and 16 ($LargeHousehold_m$). The inclusion of community-year fixed effects δ_{tc} captures any community-level effect of the UPE policy, including differences in school access or quality as well as secular and local changes in fertility over the two time periods. Variations in school-aged sibship size (number of siblings) are captured by indicators Sib_m^j, which take the value of 1 if the mother had j live children aged 0-16.[10] The matrix of controls X includes the mother's age group and schooling level, sex and age of the household head, and the wealth quintile of the household.

If the UPE policy discourages additional fertility of large households relative to smaller households, $\alpha < 0$. To be a valid measure of the causal impact of the UPE policy on fertility, a number of assumptions must hold. The first is the parallel trends assumption: absent the UPE policy, the trend in fertility change of families with more than four children should not change relative to those of smaller families. A violation of this assumption would occur, for instance, if a demographic transition were occurring in Uganda, such that household sizes were becoming smaller. In such a transition, birth rates among large households would decline faster than those of smaller households. We will provide evidence that the parallel trends assumption holds in our setting. Second, the fertility of parents with fewer than four children should not respond to the eligibility rule. A violation of this assumption could occur if parents with fewer children increase birth spacing following the policy. In that case, α is biased towards the null, and our analysis underestimates the impact of the policy on fertility.

To study whether the policy created bunching of fertility around the four-children cutoff, we estimate the following equation:

[10] Results do not change if we control for all children ever born.

$$AdditionalBirth_{mtc} = \sum_{j=1}^{10} \alpha^j PostUPE_t \times Sib_m^j + \sum_{j=1}^{10} \gamma^j Sib_m^j + X_m \beta + \delta_{tc} + \epsilon_{mtc}. \tag{2}$$

That is, we allow the response to the UPE policy to vary by the number of siblings. We expect the coefficient estimates α^j to be negative for $j \geq 4$. In addition, one could expect that there is bunching occurring at four existing children because of the discontinuous change in schooling costs at the cutoff, and that the policy response for larger households may decline with the number of existing children. This is because the per-child school fee rebate falls with the number of children. Thus, the prediction of the model is $\alpha^4 < \alpha^5 < ... < 0$.

We will also study the repeal of the eligibility policy in 2003. Our analysis will replicate the baseline regression 1 using 1995 and 2003 data. Household structure in December 2002 is used to estimate fertility from January 2003 until April 2006, when the 2003 survey was taken. The comparable fertility period for the 1995 cohort is then used in the analysis. We expect α to be ambiguously signed; however, relative to the regressions from the eligibility regime, we expect larger coefficient estimates.

Summary statistics Appendix table A1 reports summary statistics for the prereform sample (DHS 1995) and postreform sample (DHS 2000). One can see that the likelihood of additional births in the fertility period, $AdditionalBirth$, is very similar in both samples; 54 percent of women gave birth during the 34-month period. In addition, 11-13 percent of women were pregnant at the time of the interview. In terms of existing household composition, we can see that the average household had 1.5 children aged 0-16 and 1.3 children aged 0-12 in either 1996 or 1991. The UPE eligibility limit was relevant to 18-20 percent of households with four or more children in the 0-16 age category (i.e., the $LargeHousehold$ variable in regression 1). Approximately 40 percent of mothers in both samples have yet to report a birth. The table also reports other household and mother control variables. The two samples are quite balanced on those variables, with the exception of maternal education and female-headed household; both are higher in the 2000-2001 sample.

4 Results

4.1 Baseline fertility responses

Table 1 shows estimates of regression (1). Column 1 reports the simple differences-in-differences regression which only includes the $Post$ and $LargeHousehold$ indicators and their interaction. Column 2 adds individual household and mother controls but no sibship controls other than the "large household" dummy from the difference-in-differences model. Column 3 includes sibship controls and column 4 adds community fixed effects, which is our preferred specification as it properly accounts for location- and time-variant unobservable shifts in fertility. Encouragingly, estimates are very similar across specifications, including the simple difference-in-differences from column 1. All estimates are negative and range from -0.039 (column 3) to a maximum of -0.049 (column 2). Our preferred estimate from column 4 is -0.042; this indicates that families whose additional child would be ineligible for UPE reduced the likelihood of an additional birth by 4.2 percentage points. The magnitude of the impact is quite large, representing 7.8 percent of the average likelihood of subsequent birth, which was 54.2 in the 1995 sample.

Note that our results link to the UPE reform in at least two ways. First, it might have activated a standard quantity-quality trade-off response with households reducing their equilibrium household size. Since mothers with larger parities are more likely to be close to this equilibrium, we would observe a disproportionate fertility reduction among larger households. Alternatively, families are directly responding to the high marginal cost of providing education to the unborn child. Empirically, the former case should lead to a positive correlation between the fertility response and parity. In the latter, the policy should induce a large drop in fertility in households with exactly four children; this decline should be lower at higher parities. Our first piece of evidence comes from the interacted model (2); in Figure 1, we report coefficient estimates of the $PostUPE$ interacted with maternal parity after controlling for our preferred set of covariates. Coefficient estimates are close to zero and statistically insignificant for mothers with three or fewer children. We thus fail to reject the hypothesis that low parity mothers did not reduce their subsequent fertility (say, by increasing birth spacing). The estimate turns negative and statistically significant for mothers with four children; subsequent fertility did change for these mothers. Note that the coefficient for mothers with four children is not quite different from the estimate for mothers with three children (p-value of T test is 0.13). The difference between the two coefficients is equal to -0.05, which could be interpreted as the "local" treatment effect of the policy around the eligibility cutoff. The coefficient on five children is very similar but slightly noisier than the one for four children. The coefficient for six children, instead, is only slightly negative and closer to zero.[11] While we are unable to draw definitive conclusions as the estimates are not statistically distinguishable from one another, this is in line with what one could expect from bunching of the policy effect around four children.

Another way to see that the fertility response is driven by the policy cutoff is to run placebo difference-in-differences analysis assuming eligibility cutoff values that are different from four. As can be seen in Appendix Table A2, the difference-in-differences is negative but small and insignificant when the eligibility cutoff is defined as having one child or more. It becomes more negative as the cutoff moves towards four, and loses both magnitude and significance as soon as the cutoff is above four. This is what we would expect if the true cutoff was around four: as the alternative cutoff moves away from it, measurement error of the treatment variable becomes more an more severe; consistent with classical measurement error, the coefficient estimate becomes more biased toward zero.

[11] We do not report coefficients for more than six children as confidence intervals become very large due to small sample sizes.

Finally, we check the robustness of the results to alternative ways to identifying "large" households. In our baseline regression, we define the large household to be one with at least four children aged between zero and 16. Figure 2 shows point estimates of regressions with alternative definitions, starting from the most restrictive (including only children aged 0 to 10) to the most relaxed (including any child aged 0 to 18). Coefficient estimates are not at all sensitive to the definition of a large household.

Heterogeneity analysis The average treatment effects reported in the previous subsection mask a significant amount of heterogeneity of policy responses. We demonstrate this heterogeneity in Table 2, in which we report regression results for several subpopulations of interest. In particular, we report separate estimates for the urban and rural sample (columns 1 and 2 respectively), mothers with and without some secondary education (columns 3 and 4), and by three levels of household wealth (columns 5, 6, and 7). As can be seen, estimates are negative, large, and significant only for the urban, educated, and wealthy samples.

The difference in response between urban and rural areas (or rich and poor) is possibly due to a number of potential factors. It may be the case that the policy was not enforced as strongly in rural areas as in urban areas; indeed, there is some anecdoctal evidence that rural, poorer communities were less monitored by the central government (Hubbard, 2007). Alternatively, rich and urban households responded more strongly to the policy because these subpopulations were more likely to substitute more crowded public schooling with private schooling after the reform. Finally, it is possible that the lack of significance in rural areas and among the poor is driven by measurement error. With average years of schooling much lower, many older children who are counted as UPE eligible in our empirical strategy may not have attended school following the reform. Since this makes any additional child UPE eligible, the effective cutoff for rural households may be somewhat higher. Accordingly, the difference-in-differences estimate using the policy cutoff may be downward biased.

Finally, in column 8, we run a baseline specification (1) that includes the interaction between the post-UPE period and a dummy variable identifying mothers with no prior births (including births of children no longer alive). The specification allows us to separate intensive and extensive margin responses to the UPE policy, as these two might move in the opposite directions (Aaronson et al. (2014)). In the presence of extensive fertility responses, the estimate on this interaction of $PostUPE$ with no children should be positive. In our setting, we obtained a tightly estimated coefficient close to zero, indicating no extensive margin effects.[12]

Falsification test One concern with the results presented above is that the difference-in-differences estimates may be picking up secular changes in fertility. If families choose to have fewer children for reasons unrelated to the policy, we would expect to see a significantly larger reduction in fertility among high-parity women than among low-parity women over time; that is, we have a violation of the parallel trends assumption. One way to verify whether the fertility change is driven by some secular transition toward smaller families is to repeat the difference-in-differences exercise while focusing on fertility over periods of time that do not overlap with the 1997 UPE policy. Fertility transitions occur over a lengthy period of time and would result in significant and negative coefficient estimates on the difference-in-differences regressions.

In Figure 3, we carry out this exercise and report the difference-in-differences estimates when the outcome variable is having one more birth over a number of 34-month periods. We consider ten such periods, starting in 1987 for the "treated" period. All coefficients estimated over a period of time preceding 1997 are close to zero and statistically insignificant; it is only once we include fertility in 1997 or 1998 that the coefficient estimates turn sharply negative and significant. There is thus a lack of evidence against the parallel trends assumption

[12] The lack of extensive margin response cannot be fully explained by the short time horizon. We replicated the regression using a longer time horizon (as discussed in section 4.2) and found identical results (appendix table B1.)

underpinning the difference-in-differences strategy.

4.2 Long-run fertility responses

The fertility responses reported so far occur within a period of 46 months (i.e., three years and 10 months) following the onset of the policy in January 1997. We next turn to the estimation of long-run effects. This is interesting on two grounds. First, it is possible that household responses to the introduction of the UPE policy were temporary, in which case fertility could eventually catch up. In other words, a longer time period is necessary to understand if the policy led to a reduction in household size. In the online appendix, we use the 2006 DHS to show that the fertility response did not abate over a longer time horizon of 70 months (almost six years). Second, we want to demonstrate whether these responses persisted after 2003, when UPE eligibility limits were lifted.[13] Theoretically, this reduction in schooling costs should have had an ambiguous effect on fertility.

To measure the effect of this particular policy change, we take a similar approach to the one described in Section 3: using the 2006 DHS, we construct household structure at the onset of the policy change in January 2003 and measure changes in fertility between then and the onset of the survey in May 2006. We next construct a control group from the 1995 survey and replicate the baseline difference-in-differences analysis.[14] We report the result of that analysis in Table 3. After controlling for observable characteristics, the coefficient estimates are negative, highly significant, slightly larger than those in Table 1. Thus, relative to the period where schooling was uniformly expensive, fertility in the post-2003 reform period was much lower but the relaxation of the policy does not seem to have induced changes relative to the 1997-2003 period.

[13] It is unclear when, precisely, this policy was relaxed. In addition to this policy, the government also introduced secondary school reforms to encourage UPE children to continue their education.

[14] There are two differences between the specification adopted in this section and the baseline: the 2006 database does not include a wealth index, and therefore we cannot control for the wealth quintile. Additionally, the reference period is five months shorter due to differences in the timing of the 2000 and 2006 surveys.

The lack of a response to the relaxation of the UPE eligibility rules is in line with an ambiguous impact of schooling costs implied by the quantity-quality model. It may also be the result of changes in social norms–from very large families toward smaller families–that began with the 1997 reform. In the next section, we show that the period under study did experience a shift in stated preferences toward smaller families.

4.3 Desired fertility

In this section, we provide evidence consistent with a policy-induced shift in the desired fertility of women following the 1997 and 2003 reforms. For this analysis, we use the ideal number of children reported by women in 1995, 2000-2001, and 2006; the histograms for each survey round are plotted in the top panel of Figure 4. The preferred (modal) desired number of children in all waves is four, with the second most common response being six. However, the proportion preferring four children is higher in 2000 than in other periods, and fewer women in that survey round prefer having more than four children than in the other rounds. Pairwise comparisons of the three cumulative distributions using Kolgomorov-Smirnov tests indicate the distributions are indeed different (corrected p-values are all zero).

The bottom panel of the figure provides further confirmation of this pattern. We pool the three DHS years, then regress whether the woman chooses a particular number as the most desirable brood size on the DHS year, and control for the characteristics of the household. For each regression, the graph reports the coefficient estimates on the years 1995 and 2006, with the 2000 round being the excluded category. The overall pattern of the point estimates indicate a lower probability of reporting four children or less in 1995 and 2006 relative to 2000. To confirm that the distribution of desired family size temporarily shifts during the policy, in Online Appendix Table B2, we regress the number of desired children on the year dummies; the year dummies are significant. The results confirm a reduction in desired fertility in 2000, and a bounce-back by 2006.[15]

To conclude, the data strongly suggest that eligibility limits reduced desired fertility to below the eligibility cutoff, and that this response was, at least to some extent, temporary.[16] The reduction in desired fertility is consistent with the quantity-quality model. It is also consistent with the hypothesis that the policy generated a new social norm regarding household size. In particular, the fact that the government highlighted households with up to four children could have focused preferences toward that number. Given that the government did not appear to intend to control household size, this would have happened as an unexpected policy consequence. Under this hypothesis, we should also expect that, due to the new social norms around smaller family sizes, mothers' *wantedness* of the last born should have been lower among large households. However, when regressing our benchmark regression on whether the mother wanted the birth at that time (as opposed to later or never), we find insignificant DID coefficient estimates (results available).

5 Conclusion

This paper presents evidence of fertility responses to the UPE policy implemented in Uganda in 1997. Like most other UPE policies, the Ugandan reform eliminated all fees for children attending primary school for up to four children per household; the rest had to pay tuition costs as before. This policy thus created different prices of schooling for children who resided in the same community and attended the same school, but who belonged to households with different numbers and composition of children.

[15] The data includes women's reports on a coarse measure of husband's fertility preferences; in 2000-2001, 44 percent of husbands reported wanting "the same" number of children as their wives. The figure is 41 percent in 1995 and 40 percent in 2006. This is consistent with the policy influencing stated preferences for both men and women.

[16] Interestingly, this result contrasts with findings from Behrman (2015), who finds that the Ugandan UPE led to long-term declines in desired fertility. The two findings are easily reconcilable: Behrman uses the 2011 round of the DHS, which includes a significant number of women who benefited from both primary and secondary school reform (implemented in 2004). Maternal education remains an important pathway for the evolution of desired fertility in the long run, but it is absent in the short run.

We find evidence that large households (four or more children) responded to the policy by reducing subsequent pregnancies and live births. In particular, we estimate a 4.2 percentage point reduction in the likelihood of a pregnancy within the subsequent 46 months. We find that this fertility response appears only around the time of the policy, and is robust to alternative specifications. The paper provides novel evidence in support of the quantity quality model in low income countries.

Conflict of interest: The authors declare that they have no conflict of interest.

References

Aaronson, D., F. Lange, and B. Mazumder (2014). Fertility transitions along the extensive and intensive margins. *The American Economic Review 104*(11), 3701–3724.

Ainsworth, M., K. Beegle, and A. Nyamete (1996). The impact of women's schooling on fertility and contraceptive use: a study of fourteen sub-Saharan African countries. *The World Bank Economic Review 10*(1), 85–122.

Ariho, P., A. Kabagenyi, and A. Nzabona (2018). Determinants of change in fertility pattern among women in Uganda during the period 2006–2011. *Fertility Research and Practice 4*(1), 4.

Baird, S., C. McIntosh, and B. Özler (2011). Cash or condition? evidence from a cash transfer experiment. *The Quarterly Journal of Economics 126*(4), 1709–1753.

Baird, S., C. T. McIntosh, and B. Özler (2019). When the money runs out: Do cash transfers have sustained effects on human capital accumulation? *Journal of Development Economics 140*.

Becker, G. S. (1960). An economic analysis of fertility. In *Demographic and Economic Change in Developed Countries*, pp. 209–240. Columbia University Press.

Becker, G. S. and H. G. Lewis (1973). On the interaction between the quantity and quality of children. *Journal of Political Economy 81*(2, Part 2), S279–S288.

Becker, G. S. and N. Tomes (1976). Child endowments and the quantity and quality of children. *Journal of Political Economy 84*(4, Part 2), S143–S162.

Behrman, J. A. (2015). Does schooling affect women's desired fertility? Evidence from Malawi, Uganda, and Ethiopia. *Demography 52*(3), 787–809.

Bursztyn, L., A. L. González, and D. Yanagizawa-Drott (2018). Misperceived social norms: Female labor force participation in Saudi Arabia. Technical report, National Bureau of Economic Research.

Chicoine, L. (2020). Free primary education, fertility, and women's access to the labor market: Evidence from Ethiopia. *World Bank Economic Review forthcoming.*

Deininger, K. (2003). Does cost of schooling affect enrollment by the poor? Universal primary education in Uganda. *Economics of Education Review 22*, 291–305.

Doepke, M. (2015). Gary Becker on the quantity and quality of children. *Journal of Demographic Economics 81*(1), 59–66.

Fiszbein, A. and N. R. Schady (2009). *Conditional Cash Transfers: Reducing Present and Future Poverty*. World Bank Publications.

Hubbard, P. (2007). Putting the power of transparency in context: Information's role in reducing corruption in Uganda's education sector. *Available at SSRN 1100131*.

Kabagenyi, A., A. Reid, G. Rutaremwa, L. M. Atuyambe, and J. P. Ntozi (2015). Has Uganda experienced any stalled fertility transitions? Reflecting on the last four decades (1973–2011). *Fertility Research and Practice 1*(1), 14.

Keats, A. (2018). Women's schooling, fertility, and child health outcomes: Evidence from Uganda's free Primary Education Program. *Journal of Development Economics 135*, 142–159.

Kravdal, Ø. (2002). Education and fertility in sub-Saharan Africa: Individual and community effects. *Demography 39*(2), 233–250.

Lloyd, C. B., C. E. Kaufman, and P. Hewett (2000). The spread of primary schooling in sub-Saharan Africa: Implications for fertility change. *Population and Development Review 26*(3), 483–515.

MoE&S (1998). The way forward: UPE handbook: An outcome of the national conference on UPE Programme.

Munshi, K. and J. Myaux (2006). Social norms and the fertility transition. *Journal of Development Economics 80*(1), 1–38.

Museveni, Y. K. (1996, March). Y. Kaguta Museveni's speech at the launch of his election manifesto. Kampala, Uganda.

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

Qian, N. Quantity-quality and the one child policy: The positive effect of family size on school enrollment in China. In *Gender and Development.*

Rosenzweig, M. R. and K. I. Wolpin (1980). Testing the quantity-quality fertility model: The use of twins as a natural experiment. *Econometrica*, 227–240.

Rosenzweig, M. R. and J. Zhang (2009). Do population control policies induce more human capital investment? Twins, birth weight and China's "one-child" policy. *The Review of Economic Studies 76*(3), 1149–1174.

Schultz, T. P. (1997). Demand for children in low income countries. *Handbook of Population and Family Economics 1*, 349–430.

Stecklov, G., P. Winters, J. Todd, and F. Regalia (2007). Unintended effects of poverty programmes on childbearing in less developed countries: Experimental evidence from Latin America. *Population Studies 61*(2), 125–140.

Todd, P. E. and K. I. Wolpin (2006). Assessing the impact of a school subsidy program in Mexico: Using a social experiment to validate a dynamic behavioral model of child schooling and fertility. *American Economic Review 96*(5), 1384–1417.

6 Appendix: Quantity-quality model

Consider a household that derives utility u from consumption c, the number of children who survive to schooling age n, and the quality of children, q (in our analysis, the amount of education received). As in the standard model, we assume parents have a strong equitability motive; they want all children to have the same amount of education. The lifetime budget constraint is given by:

$$I = \pi_c c + \pi_q q + \pi_n n + \pi_e nq, \tag{3}$$

where π_c is the price of consumption, π_q is the cost of child quality not dependent on the number of children, π_n is the cost of the number of children independent of their quality, and π_e is the cost of providing quality q to each child. In our setting, we can think of tuition costs being reflected in the price π_e, while health and food expenditures are collected in the price π_n. In addition, if schooling reforms change the overall quality of education, we can think of those impacting π_q through a change in the returns from schooling.

Given the budget constraint (3), the first-order conditions imply shadow prices for the number of children and their quality:

$$p_n = \pi_n + \pi_e q \tag{4}$$

$$p_q = \pi_q + \pi_e n \tag{5}$$

Where $p_j = (1/\lambda)u_j(.)$ is the price for child trait j; note that the shadow price of child quality depends on the number of children n, and the shadow price of the number of children is increasing in child quality q. The marginal rate of substitution between quality and quantity is:

$$\frac{p_n}{p_q} = \frac{\pi_n + \pi_e q}{\pi_q + \pi_e n}.$$

UPE reform budget constraint Now consider a reform that reduced the cost of education from π_e to $\tilde{\pi}_e$ for up to four children. For any child above 4, households need to pay a surcharge π_d, such that:

$$\pi_d = \pi_e - \tilde{\pi}_e \tag{6}$$

In other words, the cost of schooling for the nth child is the same as in the prereform period, as long as $n > 4$. The budget set can be written as follows:

$$I = \pi_c c + \pi_q q + \pi_n n + \tilde{\pi}_e nq + D_{\{n>4\}} \pi_d q[n-4]. \tag{7}$$

Here, $D_{\{n>4\}}$ is an indicator that switches to one if the number of children is above 4.

For households for which $n \leq 4$, the constraint becomes:

$$I = \pi_c c + \pi_q q + \pi_n n + \tilde{\pi}_e nq, \tag{8}$$

while the budget set for households with $n > 4$ is:

$$I = \pi_c c + [\pi_q - 4\pi_d]q + \pi_n n + \pi_e nq. \tag{9}$$

One key difference between the two segments is that in (8) the marginal cost of an additional child is lower than before, while in (9) the marginal cost of an additional child is unchanged from before. However, in a large household ($n^* > 4$), the price of pure quality is now reduced. It is straightforward to show that the change in the marginal rate of substitution between quantity and quality of children is ambiguous when $n^* \leq 4$, whereas the MRS for $n^* > 4$ becomes:

$$\frac{p_n}{p_q} = \frac{\pi_n + \pi_e q}{\pi_q - 4\pi_d + \pi_e n}.$$

This clearly indicates that, holding income constant, the policy should induce an increase in quality and a reduction in quantity. Despite this, it is possible that the policy induces an increase in quantity if the income effect from the policy is large enough; that is, children are normal goods and the income elasticity is large.

7 Figures and Tables

Figure 1: Estimates of PostUPE $\times$ Number of existing children

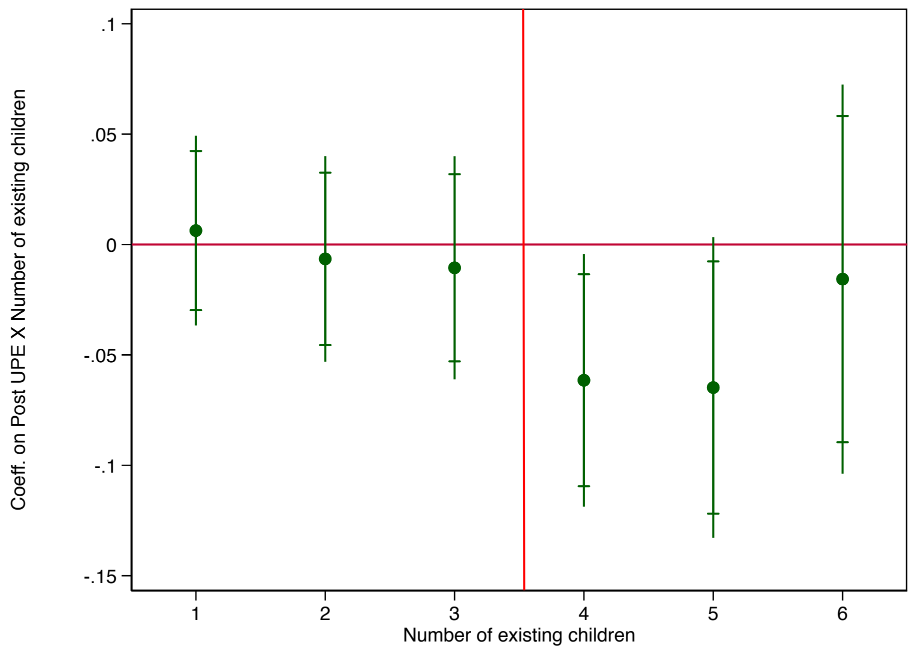

Estimates of interaction dummies between Post-UPE and the number of existing children (equation (2)). Dependent variable is having an additional birth. The excluded coefficient is the interaction with no children. Vertical lines and bars represent 90 and 95 percent confidence intervals. Estimates for 7 to 10 children omitted.

Figure 2: Difference in difference estimates, using various age groups to define large households

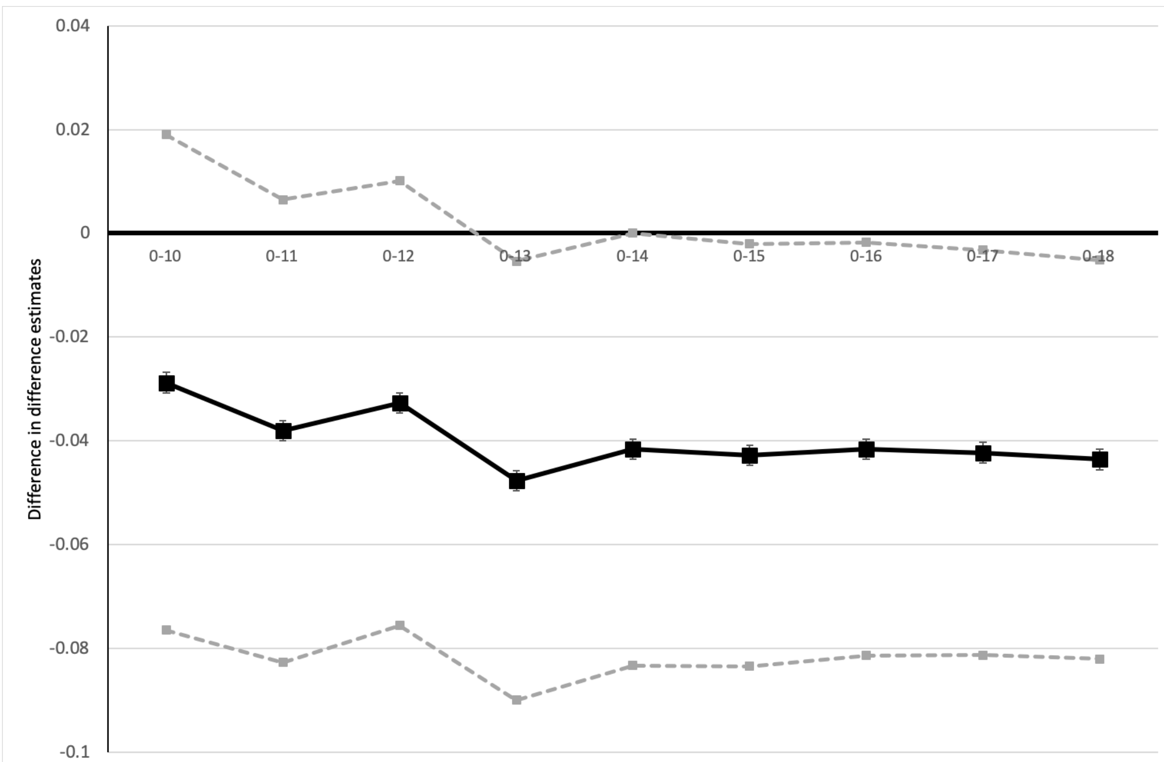

Graph of point estimates of Post UPE × LargeHousehold, where large household is defined as having four or more children within a certain age group. The age group used in the regression is defined on the x-axis. Each point in the graph is an estimate from a different regression. Regressions include the full set of controls used in table 1.

Figure 3: Placebo regressions: difference-in-differences coefficients on fertility over various 34-month periods

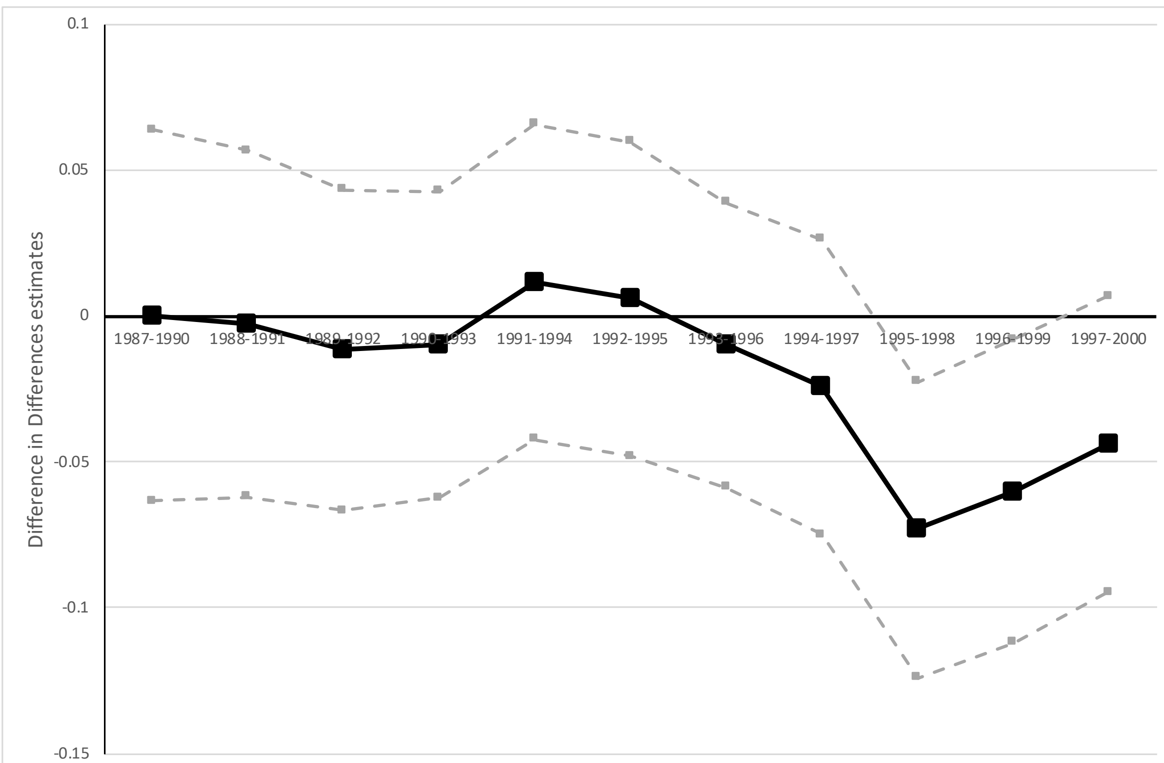

Placebo regressions of regression (1) over a 10-year period, using the DHS 2000/2001 for the treated group and DHS 1995 for the control group. The dependent variable is having an additional birth over the reference period. The reference period for treated women in each regression is indicated in the X axis. The reference period for the control group is six years before the treatment period. Sample excludes women younger than 28 years old. 95 percent confidence intervals shown.

Figure 4: Desired number of children for women before, during, and after UPE eligibility limits

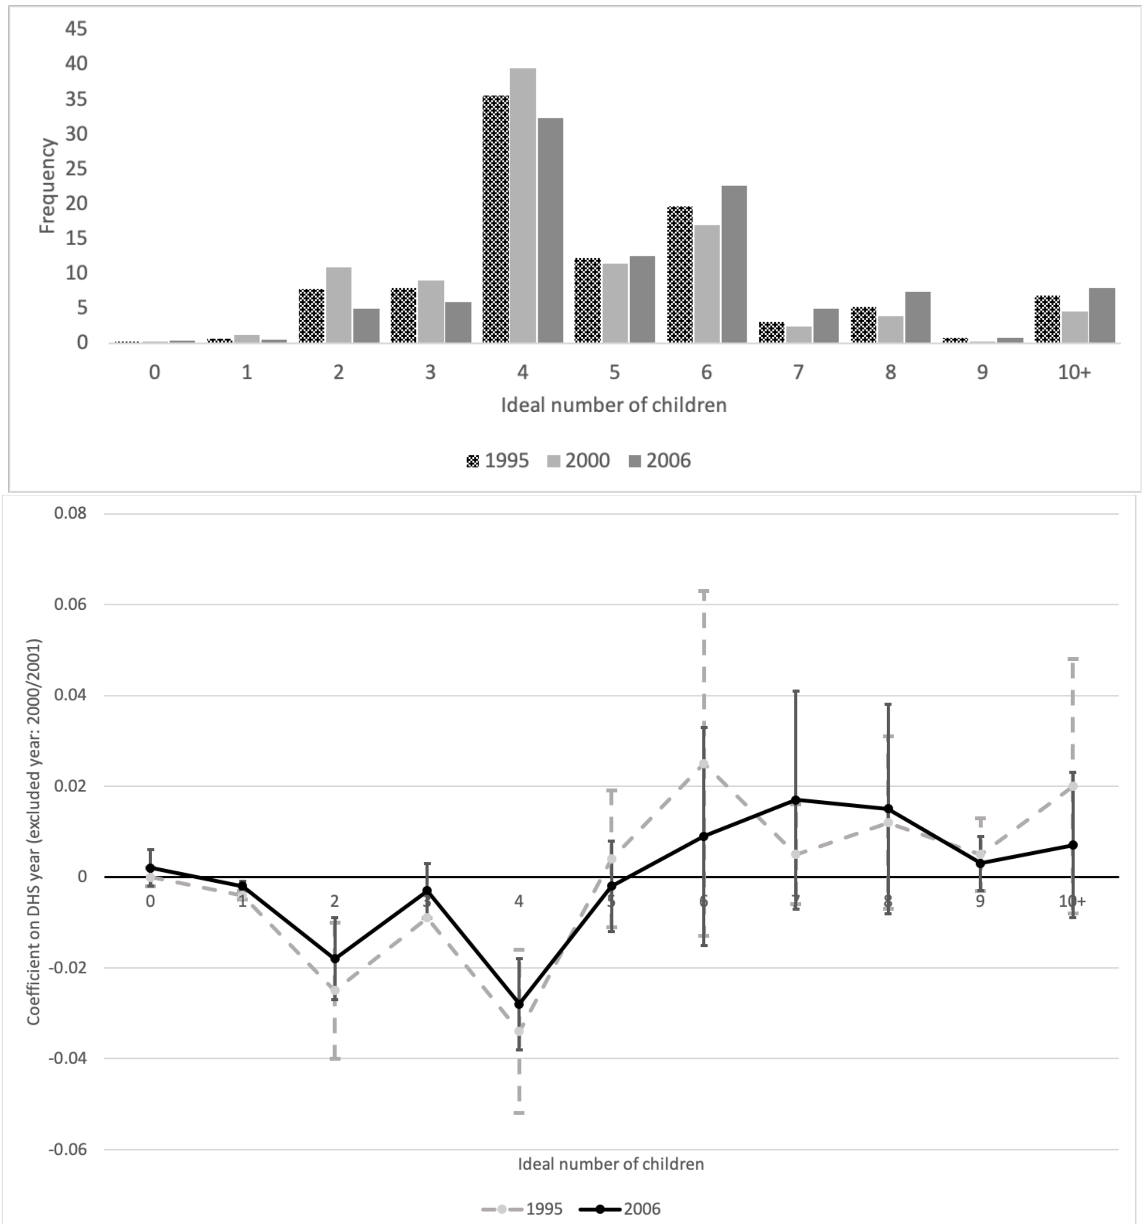

Top panel: histogram of desired number of children for mothers in three rounds of DHS. Desired fertility is topcoded at 10 children. Bottom panel: coefficient plots from separate regressions of desired number of children on DHS year. The outcome variable in each regression is whether the mother reported wanting x children, with x varying from zero to 10 or more. Regression includes dummies for each DHS round, and mother and household controls excluding siblings controls. The graph reports the coefficient estimates for the DHS years 1995 and 2006. Comparison group are mothers interviewed in 2000/2001. Robust confidence intervals illustrated by error bars.

Table 1: difference-in-differences estimates on additional births (OLS)

VARIABLES	(1)	(2)	(3)	(4)
Post UPE X Large household	-0.038*	-0.048**	-0.038*	-0.042**
	(0.021)	(0.019)	(0.019)	(0.020)
Observations	14,316	14,314	14,314	14,314
R-squared	0.009	0.281	0.302	0.337
Controls	No	Yes	Yes	Yes
Num. Siblings f.e.	No	No	Yes	Yes
Community-year f.e.	No	No	No	Yes

Regressions on women aged 15-49 using DHS 1995 and 2000-01. Dependent variable is indicator for any birth in the 46 months between January 1997 and October 2000 (for women interviewed in 2000-2001), and January 1991 and October 1994 (for women interviewed in 1995). Controls include mother 5-year age cohort, mother's education, sex of household head, household wealth quintile, rural community dummy, and year of interview. Regression 1 includes post-UPE dummy and the "big household" indicator. Number of siblings fixed effects indicate the number of live children aged 0-16 born before January 1997 (for DHS 2000-2001) and January 1991 (for DHS 1995). Errors clustered at the community level.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 2: Fertility effects across subpopulations

VARIABLES	(1) urban	(2) rural	(3) mother sec. school or more	(4) mother less than sec. school	(5) lowest wealth quintile	(6) middle wealth quintiles	(7) top wealth quintile	(8) Extensive margin
Post UPE X Large household	-0.113***	-0.017	-0.204***	-0.019	-0.000	0.004	-0.132***	-0.043**
	(0.041)	(0.023)	(0.062)	(0.022)	(0.052)	(0.026)	(0.041)	(0.021)
Post UPE X No children								-0.005
								(0.016)
Observations	4,855	9,459	3,049	11,265	2,194	7,257	4,863	14,314
R-squared	0.266	0.353	0.379	0.345	0.434	0.366	0.320	0.334
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Num. Siblings f.e.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Community-year f.e.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Regressions of Equation 1. Each regression includes a specific subsample, described in the column title. Column 8 includes all observations and an unreported indicator for having no children.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3: Long-run effects: estimates on births after UPE limits lifted in 2003

VARIABLES	(1)	(2)	(3)	(4)
Post 2003 X Big household	-0.024	-0.051***	-0.046**	-0.059***
	(0.020)	(0.018)	(0.018)	(0.019)
Observations	11,009	11,008	11,008	11,008
R-squared	0.005	0.247	0.253	0.313
Controls	No	Yes	Yes	Yes
Num. Siblings f.e.	No	No	Yes	Yes
Community-year f.e.	No	No	No	Yes

Regressions using DHS 1995 and DHS 2006 only. Dependent variable is indicator for birth in the 37 months between January 2003 and April 2006 (for post-2003 cohort), and January 1992 and April 1995 (for the pre-UPE cohort). Controls are the same as in Table 1 but exclude wealth quintile dummies (information not provided in 2006). Errors clustered at the community level.

*** p<0.01, ** p<0.05, * p<0.1

8 Appendix

Table A1: Summary statistics

	DHS 1995			DHS 2000		
VARIABLES	N	mean	sd	N	mean	sd
Fertility						
additional birth indicator	7,070	0.542	0.498	7,246	0.541	0.498
currently pregnant	7,070	0.134	0.341	7,246	0.116	0.320
number of children 0-16	7,070	1.579	1.926	7,246	1.659	1.909
large hhld (0-16)	7,070	0.185	0.388	7,246	0.198	0.398
Number of children 0-12	7,070	1.334	1.607	7,246	1.383	1.571
large hhld (0-12)	7,070	0.134	0.340	7,246	0.130	0.336
no children	7,070	0.417	0.493	7,246	0.394	0.489
ideal number of children	6,692	5.057	2.236	6,903	4.652	2.060
Hhld characteristics						
rural	7,070	0.655	0.475	7,246	0.667	0.471
female head	7,069	0.262	0.440	7,246	0.294	0.455
quintiles of wealth index	7,070	3.379	1.476	7,246	3.414	1.462
Mother characteristics						
no education	7,070	0.256	0.436	7,245	0.201	0.401
primary	7,070	0.552	0.497	7,245	0.566	0.496
secondary	7,070	0.189	0.392	7,245	0.187	0.390
higher ed	7,070	0.00354	0.0594	7,245	0.0460	0.209
age 15-19	7,070	0.230	0.421	7,246	0.233	0.423
age 20-24	7,070	0.222	0.415	7,246	0.213	0.409
age 25-29	7,070	0.187	0.390	7,246	0.183	0.387
age 30-34	7,070	0.140	0.347	7,246	0.132	0.338
age 35-39	7,070	0.105	0.307	7,246	0.108	0.310
age 40-44	7,070	0.0672	0.250	7,246	0.0755	0.264
age 45-49	7,070	0.0496	0.217	7,246	0.0560	0.230

Table A2: DID at various thresholds

Placebo threshold at child:	1	2	3	4	5	6	7	8
Post UPE X	-0.017	-0.026	-0.032*	-0.042**	-0.022	0.021	0.071	-0.032
Above threshold	(0.016)	(0.016)	(0.017)	(0.020)	(0.025)	(0.036)	(0.059)	(0.128)
Observations	14,314	14,314	14,314	14,314	14,314	14,314	14,314	14,314
R-squared	0.337	0.337	0.337	0.337	0.337	0.337	0.337	0.337
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Num. Siblings f.e.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Community-year f.e.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Dependent variable is indicator for additional birth in the reference period. Variable "Above threshold" identifies families with number of children aged 0-16 equal or larger than the threshold.

*** p<0.01, ** p<0.05, * p<0.1

Figure A1: Schematics of treatment and control periods

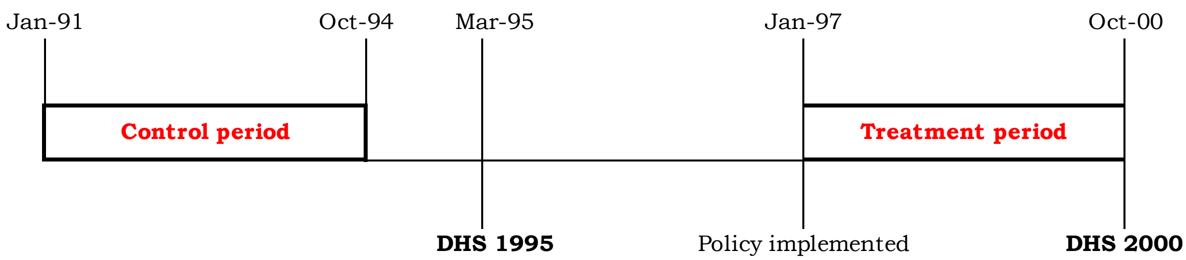

Control period uses data from the 1995 DHS while the treatment period uses data from the 2000-2001 DHS. household structure (number of live births of children aged 0-16) taken right before the two periods, in December 1990 and December 1996; additional fertility determined in the 46-month period between January 1991 and October 1994, and between January 1997 and October 2000.

9 Online Appendix

9.1 Long-run estimates of 1997 eligibility policy

In this online section, we demonstrate that the fertility response documented in the paper was not short lived. We employ two separate strategies to increase the time horizon over which fertility responses are measured and report the results in Table B1. One way to measure fertility catch-up using our baseline empirical strategy is to include currently pregnant women in our measure of $AdditionalBirth$. The inclusion of pregnant women adds 3.5 percentage points to the proportion of women who already had an additional child. Using this outcome variable, we obtain a coefficient estimate of -0.037 (column 1), which is a smaller number and thus consistent with a slight decline in the fertility response. However, we cannot reject that this estimate is different from the baseline estimate.

Our second strategy uses birth histories of women interviewed at later time periods so that a longer fertility period can be observed. Here, we construct the control group from the 2001 DHS (rather than the 1995 round) and the treated group from the 2006 DHS (rather than the 2001 round). In constructing these cohorts, we restrict the sample to mothers aged 20-49 at the time of the survey; they would have been at least 15 in either 1992 (for the control group) or 1997 (for the treatment group). We extend the fertility period under analysis by an additional 24 months (two years), giving us a response period of almost six years. The length of the evaluation period for the control group ends in October 1996, that is, right before the UPE policy was implemented. The fertility window for the treated group goes from January 1997 to October 2002. This period also ends right before the 2003 reform, which extended free tuition to all children regardless of household size. It is thus clear that our evaluation period cannot extend any further.

We first demonstrate the comparability of this alternative sample by replicating our baseline result from Table 1 (column 2). We obtain a coefficient estimate (-0.082) that is as precisely estimated but twice as negative as the -0.042 reported in column 4 in Table 1. The estimate is within the 95 percent confidence interval of the -0.042 estimate from Table 1.[17]

Having established the comparability of the two samples, we next estimate (1) on the longer fertility period (70 months) and report the result in column 3. The coefficient estimate falls slightly from -0.082 to -0.072 but remains highly significant. This is strongly suggestive that fertility effects are persistent. Finally, given the longer time period, it is possible to explore whether ineligible cohorts have also reduced the number of additional births, conditional on having had at least one more child. In column 4, we restrict the sample to mothers who gave birth at least once in the fertility period, and run the difference-in-differences regression on the number of births. The coefficient estimate is negative but small and noisy, indicating the main impact is on having one more birth and not on having a second one.

[17] One possible explanation for this is that this alternative sample excludes women aged 15-19. Reestimating (1) on the 1995 and 2000-2001 sample while excluding the 15- to 19-year-old age group does indeed lower the estimate from -0.042 to -0.053.

Table B1: Long-run effects: estimates over longer reference periods

VARIABLES	(1) Incl. pregnant women	(2) Births to 2000	(3) Data: DHS 2000-01 and 2006 Births to 2003	(4) Extensive margin	(5) Num. of births
Post UPE X Large hhld	-0.036*	-0.068***	-0.068***	-0.068***	-0.054
	(0.020)	(0.021)	(0.022)	(0.022)	(0.051)
Post UPE X No children				-0.002	
				(0.019)	
Observations	14,314	11,142	11,142	11,142	8,405
R-squared	0.328	0.243	0.243	0.243	0.214
Controls	Yes	Yes	Yes	Yes	Yes
Sibship size f.e.	Yes	Yes	Yes	Yes	Yes
Community-year f.e.	Yes	Yes	Yes	Yes	Yes

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Column 1 regresses equation 1 on an indicator for women who have had an additional birth during the reference period or were pregnant at the time of the survey. Columns 2-4 use the alternative DHS 2000-2001 and 2006 data. The dependent variable in column 2 and 3 is an indicator for an additional birth in the 70-month reference period. Column 4 further restricts the sample to women who experienced at least one birth during the 70-month reference period. The outcome variable is a count of the number of births.

*** p<0.01, ** p<0.05, * p<0.1

Table B2: Desired fertility across cohorts

VARIABLES	(1)	(2)	(3)
2000/01 DHS	-0.390***	-0.333***	-0.356***
	(0.073)	(0.046)	(0.113)
2006 DHS	0.348***	-0.113**	-0.112
	(0.074)	(0.057)	(0.087)
Observations	19,375	19,373	19,373
R-squared	0.022	0.225	0.280
Controls	No	Yes	Yes
Sibship-size f.e.	No	No	No
Community-year f.e.	No	No	Yes
P-value 2000 = 2006	0	4.29e-05	0.00135