

The Consequences of Child Labor:

Evidence from Longitudinal Data in Rural Tanzania

Kathleen Beegle
Rajeev H. Dehejia
Roberta Gatti
Sofya Krutikova

The World Bank
Development Research Group
Macroeconomics and Growth Team
July 2008



Abstract

This paper exploits a unique longitudinal data set from Tanzania to examine the consequences of child labor on education, employment choices, and marital status over a 10-year horizon. Shocks to crop production and rainfall are used as instrumental variables for child labor. For boys, the findings show that a one-standard-deviation (5.7 hour) increase in child labor leads 10 years later to a loss of approximately one year of schooling and to a substantial increase in the likelihood of farming and of marrying at a younger age. Strikingly, there are no

significant effects on education for girls, but there is a significant increase in the likelihood of marrying young. The findings also show that crop shocks lead to an increase in agricultural work for boys and instead lead to an increase in chore hours for girls. The results are consistent with education being a lower priority for girls and/or with chores causing less disruption for education than agricultural work. The increased chore hours could also account for the results on marriage for girls.

This paper—a product of the Growth and the Macroeconomics Team, Development Research Group—is part of a larger effort in the department to study the consequences of child labor. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The author may be contacted at kbeegle@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

**The Consequences of Child Labor:
Evidence from Longitudinal Data in Rural Tanzania***

Kathleen Beegle
The World Bank

Rajeev H. Dehejia
Tufts University
and
NBER

Roberta Gatti
The World Bank
and
CEPR

Sofya Krutikova
University of Oxford

* We thank seminar participants at the CSAE Conference at Oxford University, the Child Labor Conference at Indiana University, CIDE, The Hebrew University, IDEI at the University of Toulouse, the IZA conference on Migration, the Paris School of Economics, Rutgers University, Tel Aviv University, and the conference on child labor organized by UCW and Centre d'Economie de la Sorbonne for helpful comments. All errors are our own. Contact information: kbeegle@worldbank.org, rajeev@dehejia.net, rgatti@worldbank.org, sofya.krutikova@economics.ox.ac.uk.

I. Introduction

This paper exploits a unique longitudinal data set from Tanzania to examine the impact of child labor on education, employment choices, and marital status over a 10-year horizon. The question is important for many reasons. The assumption that work is harmful to children's development underpins both the theoretical literature and the policy debate on child labor. For example, from the policy perspective, there is a general perception that the worldwide returns to eliminating child labor are substantial (see International Labour Organization [ILO], 2003). However, the evidence that rigorously quantifies the consequences of child labor is limited.

Theoretically, it is unclear whether and to what extent child labor is harmful. In rural settings in developing countries (which is where more than 70 percent of the world's child labor occurs; ILO, 2002), child labor tends to be a moderate-intensity activity, while at the same time schooling is part time and intermittent. Child labor may also provide the child with work experience that subsequently could be rewarded in the labor market. The overall effect is ambiguous.

Empirically, although there is a large and growing body of evidence that child labor is harmful (reviewed in Section 2), the existing literature has many limitations, some of which we seek to overcome in this paper. First, most papers focus exclusively on schooling as an outcome. Although schooling is central, it is also important to consider other outcomes, in particular outcomes that allow us to measure possible effects of child labor on economic activity -- in this paper we look at occupation, crop choices, and labor productivity. Second, most papers in this literature examine correlations, rather than causal relationships. There are many reasons in the context of child labor why correlation need not imply causation. Both between and within households there is selection along observable (education, wealth, occupation, children's age,

and gender) and unobservable (social networks, concern for children, child ability) dimensions. In this paper we use an instrumental variables strategy which we describe below. Third, almost all papers in this literature are confined to looking at the contemporaneous or short-horizon effect of child labor. While there is interesting evidence from Brazil, an upper-middle income country, using retrospective rather than longitudinal data on child labor (Emerson and Souza, 2007), there is no direct evidence on the longer-run impact of child labor in low-income countries.

Our strategy is to examine the long-run consequences of child labor using the Kagera Health and Development Survey in Tanzania, which consists of five waves, spanning 13 years. We study the relationship between child labor (measured by hours spent on household chores and in economic activity) in early waves (1991-1994) and outcomes (such as education, whether or not the individual is farming, choice of cash versus subsistence crop, marital status, and labor productivity) in the final wave (2004). We instrument for child labor using crop and rainfall shocks. In previous work (Beegle, Dehejia, and Gatti, 2006), we have shown that crop shocks are predictive of child labor; we argue below that they, along with rainfall shocks, are plausible instrumental variables (i.e., that they satisfy conditional exogeneity and the exclusion restriction).

We find that child labor is causally associated with reduced educational attainment (as measured by years of schooling and by an indicator for completion of primary school). Interestingly, this result appears to be entirely driven by the sample of boys, for whom a one standard deviation (5.7 hour) increase in child labor implies nearly one year less of schooling. Boys who worked when young are more likely to be farming (as opposed to earning a wage). We do not find that child labor is associated with discernible differences in the choice of crop (cash

versus subsistence) or subsequent migration. For girls, the only robustly significant effect of child labor is an increased probability of being married 10 to 13 years later.

The paper is organized as follows. Section II briefly reviews the existing literature. Section III introduces the data, and Section IV describes the empirical methodology and discusses in detail the plausibility of our instrumental variables approach. Results are presented in Section V, and Section VI discusses refinements and extensions. Section VII concludes.

II. Literature Review

There is a rapidly expanding literature on child labor. In this section we briefly review key findings that put our own work in perspective. Edmonds (2007) offers a more detailed survey.

A large literature has established a negative correlation between child labor and school attainment. For example, Patrinos and Psacharopoulos (1995) show that factors that predict an increase in child labor also predict reduced attendance and an increased chance of grade repetition; Patrinos and Psacharopoulos (1997) further show that child work is a significant negative predictor of age-grade distortion. A number of papers have used test scores as an outcome. These include Akabayashi and Psacharopoulos (1999), who present evidence that children's reading competence (as assessed by their parents) decreases with child labor hours, and Heady (2003), who finds a negative relationship between child labor and objective measures of reading and mathematics ability in Ghana.

A more recent literature tries to estimate causal effects rather than correlations. These papers use a number of strategies. Boozer and Suri (2001) use regional variation in rainfall as a source of exogenous variation in child labor, and find that a one hour increase in child labor

leads to a 0.38 hour decrease in contemporaneous schooling. Cavalieri (2002) uses propensity score matching and finds that child labor is associated with a 10 percent reduction in the probability of being promoted to the next grade.

Papers using an instrumental variables strategy include Ray and Lancaster (2004), Beegle, Dehejia, and Gatti (2005), and Bezerra, Kassouf, and Arends-Kuenning (2007).¹ Each of these papers has strengths and weaknesses. Ray and Lancaster use micro data from seven countries, but their instruments (household measures of income and assets, and water, telephone, and electricity infrastructure) are unlikely to satisfy the exclusion restriction.² Beegle, Dehejia, and Gatti (2005) use community rice price as an instrument for child labor in Vietnam.³ They estimate that child labor reduces the probability of being in school by 30 percent and educational attainment by 6 percent, but are limited to looking at outcomes over a 5-year horizon. Bezerra, Kassouf, and Arends-Kuenning (2007) use city population, state-level schooling, and literacy rates to instrument for child labor in Brazil. They find that working seven hours or more per day results in a 10 percent decrease in test scores relative to students who do not work. However, their instruments are likely to be correlated to city and state unobservables, and as such are unlikely to satisfy the exclusion restriction. Finally, Ravallion and Wodon (2000) use between-village variation induced by a food-for-school program in Bangladesh; they find that the program led to a significant increase in schooling, but only one-eighth to one-quarter of the increase in schooling hours is explained by decreased child labor.

¹ Krutikova (2006) uses the same data as this paper and focuses on educational attainment.

² For instance, they find that in Belize the initial hour of child labor leads to a reduction in years of schooling by 2.6 years. Note that in some cases they find the marginal impact of child labor to be positive. In particular, for Sri Lanka, the impact is positive for all school outcomes.

³ Beegle, Dehejia, and Gatti (2005) also examine health outcomes for children, as do O'Donnell, Rosati and Van Doorslaer (2005) also in Vietnam.

The literature examining the link between child labor and subsequent labor market outcomes is much more limited. Indeed, the only papers we are aware of are Beegle, Dehejia, and Gatti (2005) and Emerson and Souza (2007). Beegle, Dehejia, and Gatti (2005) find that child labor is associated with a significant increase in wages 5 years later; the wage increase is sufficiently large to offset the cost of displaced education. Emerson and Souza (2007) instrument for child labor and child schooling using the number of schools per child in the state, the number of teachers per school, and GDP per capita at age 12. They find that, even controlling for completed schooling, child labor has a negative effect on adult earnings. The strength of this paper is a large, nationally representative data set, but its limitations are that child labor is measured retrospectively (in adulthood) and only for the sample of adults who are working for a wage. In contrast with these papers, the present paper observes child labor as it occurs and is able to follow individuals over a 10-year horizon.

III. Data Description

III.1 Data set

The Kagera Region of Tanzania is located on the western shore of Lake Victoria, bordering Uganda to the north and Rwanda and Burundi to the west. The population (1.3 million in 1988, about 2 million in 2004) is overwhelmingly rural and primarily engaged in producing bananas and coffee in the north and rain-fed annual crops (maize, sorghum, cotton) in the south. This study uses baseline data from the Kagera Health and Development Survey (KHDS), a longitudinal socioeconomic survey conducted from September 1991 to January 1994 covering the entire Kagera region (World Bank, 2004). Because adult mortality of the working age population (15-50) is a relatively rare event and HIV/AIDS was unevenly distributed in Kagera,

the KHDS household sample was stratified. In order to capture a higher percentage of households with a death while retaining a control group of households without a death, stratification was based on agro-climatic features of the region, levels of adult mortality from the 1988 Census (including both high and low mortality areas), and household-level indicators thought to be predictive of elevated adult illness or mortality.

In 2004, another round of data collection was completed (Beegle, De Weerdt, and Dercon, 2006a). The goal of the KHDS 2004 was to re-interview the sample of 6,210 respondents from the 1991-1994 survey; this excludes 169 individuals who died over the course of the baseline rounds. In addition to the household survey, the KHDS 2004 included additional community-level surveys consistent with those carried out in the 1991-1994 rounds. A community questionnaire was administered to collect data on the physical, economic and social infrastructure of the baseline communities, as well as shocks experienced at the community level. Over the course of 10-13 years, it was anticipated that a substantial number of individuals would have migrated from the dwelling occupied in 1991-1994. Considerable effort was made to track surviving respondents to their current location, be it in the same village, a nearby village, within the region, or even outside the region.

Because of the long time frame of the KHDS panel, we are able to study the behavior of children in conjunction with outcomes for these children as young adults. Among children ages 7-15 studied in Beegle, Dehejia, and Gatti (2006), 75 percent were re-interviewed in 2004, 21 percent were not located, and 4 percent were deceased. Among the children we study here (for details on the sample restriction see Section IV.1), 76 percent were re-interviewed in 2004. Of these, 18 percent had moved far from their original village but still resided in Kagera, 11 percent resided outside Kagera but in Tanzania, and 2 percent were residing in Uganda. These children

were, on average, 11 years old in their last interview from the baseline rounds. By 2004, they were almost 23 years old (Table 1).

III.2 Descriptive statistics

Our definition of child labor is the total hours spent working in economic activities and chores in the previous week (including fetching water and firewood, preparing meals, and cleaning the house). Economic activities for children consist predominately of farming, including tending crops in the field, processing crops, and tending livestock. We include chores as well as economic activities because the concept of child labor typically (e.g., in the ILO standard) encompasses both,⁴ although we will distinguish between the two to explain differences in our results for boys and girls. Children in the sample work on average 16.8 hours per week, of which 10 are spent on chores (Table 1). Girls spend on average 2.5 hours more than boys working on household chores; this difference is more pronounced among older girls. More than 90 percent of children have worked at least one hour in one of the baseline waves.

We use two instrumental variables: household crop shocks and rainfall shocks. Household crop shocks are measured as an indicator variable of crop lost to pests and fire during the baseline interview period. Rainfall was measured at 21 weather stations from 1980 onward. For each household, we use rainfall data from the nearest weather station. We construct the rainfall shock as the deviation of the total rainfall in the short and long rainy seasons preceding the interview from its 25-year average, scaled by its standard deviation. Table 1 shows that the average proportion of crop loss is 0.34. The mean of the rainfall shock is 0.21 (i.e., one-fifth of a

⁴It should also be mentioned that the concept of child labor does not necessarily refer to simply any work done by a child, but, rather, work that stunts or limits the child's development or puts the child at risk. However, in survey data it is difficult to isolate the portion of time spent working on the farm that qualifies under this nuanced definition.

standard deviation more rain than the weather-station specific norm).⁵ We discuss the plausibility of these instruments in Section IV.2 below.

Our education outcome variables are years of schooling and an indicator variable for having completed seven or more years of education (primary level). Individuals in the sample have an average of 6.4 years of schooling and 78 percent have completed primary school. We measure labor market outcomes with a range of variables, including whether the individual earns a salary or is farming and among those farming whether the individual is growing cash crops. As the economy in the Kagera region is based mainly on extensive farming, whether the individual earns a salary or is involved in cash cropping (mainly tobacco and coffee, rather than subsistence farming) are important indicators of success. We also examine the probability of migrating from the village⁶ and the (imputed) marginal productivity of labor in agriculture.⁷ The literature has suggested that plot-specific experience could be an important element of the rural economy (Rosenzweig and Wolpin, 1985). It is possible that child labor contributes to plot-specific experience, in which case we would see child labor associated with an increased likelihood of farming, lower individual mobility, or higher labor productivity.

Finally, we explore whether child labor significantly affects marital status. This is particularly interesting for our sample of girls, who tend to work more hours than boys, especially on household chores.⁸ Since marriage is universal in Tanzania, we are examining the influence of child labor on the likelihood of earlier marriage. Age at marriage has been shown to

⁵ By construction, the rainfall deviation variable has a mean of zero and a standard deviation of one over the entire sample period, though not necessarily within sub-periods.

⁶ In wave 5, 70 percent of re-interviewed individuals in the sample were still living in the same or in neighboring villages. Physical mobility is associated with significantly higher income gains for panel respondents (Beegle, De Weerd, and Dercon, 2006b).

⁷ The marginal productivity of labor is imputed by estimating an agricultural production function. The value of total agricultural output is regressed on all inputs – including labor and implements – that are used in production. The procedure is based on Jacoby (1993).

⁸ For example, girls between 10 and 15 work 22 hours per week (15 of which are spent on household chores), as opposed to 18 hours for boys (11 of which are spent on household chores).

be associated with worse outcomes for women and their children, including increased health risks as well as potentially “worse” marriage matches.⁹

IV. Empirical Methodology

IV.1 Specification

We are interested in the relationship between outcomes in wave 5 (including education, occupation, and marital status) and the level of child labor intensity (which we measure through mean child labor hours in waves 1 to 4). An OLS regression of the form

$$Y_{i,t} = \alpha + \beta T_{i,t-10} + \gamma X_{i,t} + \varepsilon_{i,t}, \quad (1)$$

where $Y_{i,t}$ are outcomes in wave 5, $T_{i,t-10}$ is mean child labor hours in waves 1 to 4, and $X_{i,t}$ are household and community-level controls, is most probably not suitable for estimating a causal relationship. The principal concern is omitted variable bias. The child labor decision is likely to be correlated with both household- and child-level covariates, not all of which will be observable to the researcher. For example, though we can control for parents’ education we cannot control for their discount rates. At the child level, we have few covariates other than age and sex, and thus, for example, cannot control for ability. Reverse causality is less of a concern because the outcome is measured 10 years after child labor intensity.

We address concerns with the OLS specification using an instrumental variables strategy. Our instruments, $S_{i,t-10}$, are indicators of household agricultural shocks and rainfall shocks in waves 1 to 4. Thus our basic specification is a two-stage least squares procedure of the form:

$$T_{i,t} = a + bS_{i,t} + cX_{i,t} + v_{i,t} \quad (2)$$

⁹ Younger mothers are more likely to suffer from micronutrient deficiencies and be unaware of the health risks associated with pregnancy; they are also more likely to have children soon after marriage with increased risk of maternal and infant mortality (World Bank, 2007). Younger ages at marriage may result in curtailed education for girls, although it is difficult to ascertain the causality. In any case, a younger bride may be less able to assert power and authority in her marriage especially given that women marry men who are on average several years older.

$$Y_{i,t} = \alpha + \beta \hat{T}_{i,t-10} + \gamma X_{i,t} + \varepsilon_{i,t}. \quad (3)$$

In practice, we also include interactions of one of the instruments (crop shocks) with a range of baseline characteristics (including age and sex) to increase predictive power in the first stage.

We impose several restrictions on the sample we examine. Following our previous work, we consider children between the ages of 7 and 15 in the baseline survey. Note that the prevalence of work among younger children is low. Likewise, by most definitions, working at age 16 and above would not be viewed as a particularly serious form of child labor. We also have information on whether children have ever been to school by wave 4. Tabulation of this variable shows that only 32 percent of 7-8 year olds had attended school at some point in time, which is consistent with the widespread tendency to delay enrollment, while among children age 13 and above only 12 percent had never been to school. It is unlikely that these older children who have not yet enrolled will enroll in the future. At the same time, the data suggest that, in response to a shock, households are more likely to employ the labor of the older, more productive children. Because of this, if we include these children in our sample, we would be likely to find a strong negative correlation between years of schooling and child labor. As a result, our main sample includes all 7-15 year olds who were in school at the relevant wave and those children who had not yet entered school but were still young enough to have a chance to enroll (7-9 year olds). We also present results for the full sample of 7-15 year old children.

IV.2 First-stage results and the plausibility of the instruments

In this section, we discuss whether crop and rainfall shocks plausibly satisfy the requirements for a valid instrumental variable.

Relevance

Both crop shocks and rainfall shocks affect the agricultural production function and as such could have an effect on the use of child labor. Crop shocks reflect occasions of agricultural stress or crisis, which are moments when the incremental value of child labor may be particularly high. The impact of variation in rainfall is more ambiguous. Plentiful rainfall could be a positive shock to productivity, and, depending on the agricultural production function, could increase the marginal productivity of child labor. At the same time, extreme outliers of the rainfall deviation variable – essentially floods – are also possible, and their impact on the use of child labor is less clear. As it turns out, we do not observe any floods in the first four waves of the survey.

Our previous work and the estimates presented here confirm that crop shocks are significant predictors of child labor. Table 2 reports estimates from a first-stage regression where total child labor hours are regressed on an indicator of crop shocks, rainfall shocks, and other regressors which include region, parental education, log per capita expenditure, and household size (column (1)). The preferred first stage specification also includes crop shocks interacted with age, a female dummy, and a three-way interaction between crop shock, age and female (column (2)).

The direct effect of a crop shock is a half hour increase in the number of hours worked per week. The positive association between rainfall deviation and child labor, which is only marginally significant (at 13 percent), suggests that above average rainfall increases child labor. The effects of the crop and rainfall shocks differ significantly by both age and sex; our preferred first-stage specification, therefore, includes interactions between crop shocks, age, and sex. These crop shock terms along with the rainfall deviations are jointly significant at the 1 percent level, with an F-statistic of 8.6 (column (2)). Crop shocks have a larger effect on child labor

among older children. While they are associated with a decrease in the amount worked by younger children, the magnitude of their positive effect on the amount worked by older children is greater. The association between rainfall deviations and child labor remains positive, irrespective of the age and sex of the children. Exploring the relationship between shocks and child labor further, we split the results by sex and type of work (chores and economic activities) in columns (3) to (6). We find that the incremental effect of the crop shock impacts girls' labor hours in both economic activities and chores, while among boys it increases economic activities to a greater degree than chores. Finally, it is worth noting that since younger children are over-represented in the sample used in this paper¹⁰ the overall positive effect of the crop shocks on child labor is smaller than it would be if all individuals between the ages of 7 and 15 were included in the sample.

Exogeneity

Table 3 examines the plausibility of the exogeneity requirement for our instruments. In Columns (1) to (2) crop shocks are regressed on household characteristics and lagged crop shocks. Column (1) shows that the occurrence of crop shocks is uncorrelated with such household characteristics as household size and parental education. Column (3) shows that this is also mostly true for the rainfall shocks, with the exception of a positive correlation with mother's secondary education. It is worth noting that less than 2 percent of the children in the sample have mothers with more than a primary education and that we control for parents' education in the second-stage specification.

¹⁰ As mentioned before, in order to avoid positive bias in our results, we omit children over the age of 10 who had not attended school by the time of the baseline survey.

Column (2) examines the correlation of crop shocks over time. We find a pattern of decaying correlations: there is a 10 percent correlation between shocks in adjacent survey rounds, but the correlation decreases to three percent for shocks more than a year ago. None of these effects is statistically significant. The fact that the occurrence of shocks is not significantly correlated over time suggests that the likelihood of being hit by a shock is not driven by time-invariant unobservable household characteristics: i.e., if shocks were driven by household unobservables, then we would expect to see that a household that experiences a shock in round 1 will also be more likely to be shocked in rounds 2, 3, and 4, which we do not find.

Overall, these results suggest that crop and rainfall shocks are plausibly uncorrelated with respect to household characteristics, and crop shocks are uncorrelated over time.¹¹ The one exception to this – mothers with 7 or more years of education – can be (and is) readily included as a control in the second stage.

The Exclusion Restriction

The remaining concern is whether the instruments satisfy the exclusion restriction, i.e. that they affect education and labor outcomes only through child labor. The relevance of this concern is supported by an influential strand of literature suggesting that transitory shocks can have long-term consequences for households (see, for example, Ravallion and Lokshin, 2005). We investigate this concern in a number of ways. First, we examine the reduced form effect of agricultural and rainfall shocks in waves 1 to 4 on a range of outcomes 10 to 13 years later.¹² In particular, we regress wave 5 measures of household wealth, including values of physical and

¹¹ By construction, rainfall deviations are correlated over time since they are calculated as the annual deviations from a 25 year average.

¹² This could be extended by verifying that shocks are not correlated with such causes of attrition in the sample as mortality and destitution; this is work in progress.

business assets, durables, farm equipment, and land on the average level of crop and rainfall shocks in waves 1 to 4. The effects are uniformly insignificant (Table 4). These results suggest, albeit indirectly, that given initial conditions, shocks in waves 1 to 4 did not have permanent effects on wealth variables in wave 5. In other words, they support the hypothesis that while shocks account for significant variation in child labor, their effects on other variables are likely to be short term.

Second, we exploit cross-sectional variation in the size of shocks to test whether our second stage results are driven only by large shocks or whether we obtain similar results for smaller shocks. Smaller shocks are a priori less likely to have long-lasting impacts on households, except through their impact on contemporaneous variables (child labor). We present these results below (in Section VI.3, along with other robustness checks).

Finally, we use the adults in our data as a comparison group for children. In particular, using the same specification that we use for children, we test whether labor hours for adults (those aged 20 and older) in waves 1 to 4 have a significant impact on outcomes in wave 5 when instrumented with agricultural shocks. Adults do not, by definition, participate in child labor and their outcomes in wave 5 should, therefore, not be affected by instrumented labor in waves 1 to 4, unless either the exclusion restriction is violated and the instruments affect the outcomes directly through some channel other than child labor, or the instruments also have an impact on adult labor which in turn affects outcomes. We present and discuss these results in Section VI.2 below.

V. Results

V.1 Baseline OLS and reduced-form specifications

Prior to discussing our instrumental variables results, we present the results of OLS and reduced-form specifications.

Table 5 presents results from an OLS regression of our outcomes as measured in wave 5 on average child labor over waves 1 to 4. Although the results are not statistically significant and are also likely to be biased, they are useful baseline estimates for comparison to our instrumental variables results. Child labor is associated with reduced schooling, which is measured by number of years of schooling attained and the probability of completing primary education. The negative association also holds for the probability of staying in or near the village (i.e., staying in the same village in wave 5 as in the last round of the baseline) and the probability of being a waged worker in the last 12 months. The results further suggest that child labor is positively correlated with the probability of being a farmer in the last 12 months, growing cash crops, being married, and labor productivity. Splitting the sample by gender shows that among girls child labor is associated with a statistically significant increase in the probability of marriage (which, given the sample characteristics, typically implies marrying at a younger age). However, because of the potential sources of bias in this specification, as discussed in Section IV.1, we do not interpret these coefficients causally.

V.2 *Instrumental variables estimation*

We now discuss our instrumental variables estimates of the impact of child labor hours on the outcomes of interest. In the first stage, labor hours are predicted from a regression of child labor hours on shocks and their interactions (Table 2 and equation (3) above).¹³

Table 6, columns (1) and (2), show that the instrumental variables estimates of the effect of child labor on education are negative and statistically significant. A one standard deviation increase in child labor hours (5.7 hours) is associated with a decrease of nearly half a year of schooling and an 8.8 percentage point reduction in the chance of completing primary school. These results are in line with those obtained for Vietnam by Beegle, Dehejia, and Gatti (2005). Both papers also find that IV effects are greater than OLS effects.

The direction of the bias in the OLS estimates sheds some lights on how families make decisions regarding which children in the household work more. To the extent that families send the least gifted children to work and skills in the classroom and in the field are positively correlated, OLS would overestimate the impact of child labor on schooling, relative to the causal effect (as estimated by IV).¹⁴ However, our results instead lend support to the view that families send their most gifted children to work. Parents may decide to use their most talented children if children's work in response to a shock is sufficiently important. Another possibility is that there is significant attenuation bias in the result due to measurement error. This is not implausible: our measure of child labor (hours worked in the week prior to the survey) is likely to be very noisy.

¹³ Note that as we have on average 2.3 children in the relevant age range per household, the second stage could be estimated with household fixed effects. However, there is limited variation in education outcomes across children within a family so a second stage fixed effects specification is unable to estimate the coefficients of interest with any degree precision.

¹⁴ This validates one of the key predictions of the model presented in Horowitz and Wang (2004).

In column (3), we see that child labor does not appear to be significantly associated with migration.¹⁵ In contrast, column (4) shows that boys (although not girls) who worked when young are significantly more likely to be farming in adulthood; a one standard deviation increase in child labor results in an 18 percentage point increase in the likelihood of farming in adulthood. The choice between farming cash or subsistence crops is, however, unaffected by working in childhood (column (5)).¹⁶ Similarly, whether the individual had a wage or salary job in the past 12 months is also not explained by working in childhood (column (6)). It is interesting to note that child labor is associated with a significant reduction in the marginal productivity of labor in agriculture (column (8)). A one standard deviation increase in child labor leads to a 3 percent reduction in the marginal productivity of labor in agriculture; this effect is significant at the 5 percent level.

Our result on the increased likelihood of farming can be rationalized in the Rosenzweig and Wolpin (1985) framework in which child labor imparts plot-specific experience that – as opposed to formal education – is difficult to transfer to other activities. An individual who works as a child thus benefits from locking himself into farming rather than seeking opportunities in other sectors or by migrating. Of course, one would have to believe that such plot-specific knowledge is not being picked up in our estimates of the marginal productivity of labor in agriculture, which goes down rather than up. Alternatively, our results on farming and the marginal productivity of labor could simply be viewed as implications of reduced education, which can both narrow opportunities outside the agricultural sector and – if education and labor are complementary in the production function – reduce the marginal productivity of labor.

¹⁵ Results do not change if we include in the sample the children who could not be recontacted between baseline and wave 5.

¹⁶ Note that farming and working for a wage are not mutually exclusive. Of 1,311 people: 180 do neither; 142 working for wage/salary, not farming; 648 farming, not working for wage/salary; 341 both farming and working for wage/salary.

Finally, in column (7), we find that child labor is associated with a significant increase in the probability of marriage. A one standard deviation (5.7 hour) increase in child labor leads to a 16 percentage point increase in the probability of marriage by wave 5. As noted in the discussion above, since marriage is almost universal in Kagera, this result suggests that child labor is associated with earlier marriage. The positive effect of working in childhood on the probability of earlier marriage is particularly strong for boys; a one standard deviation increase in child labor increases the probability of a boy being married by wave 5 by over 30 percentage points. This result may reflect the increased value on the marriage market of boys who have more agricultural experience or be the byproduct of reduced educational opportunities.

Strikingly, splitting the results by gender reveals that the relationship between child labor and education is driven by boys. One possible explanation of this finding is that girls may be involved in forms of child labor that are less harmful to education than the work that boys do. Splitting the first stage results into chores and economic activities reveals that while the incremental effect of a crop shock for girls is an increase in chore hours and a reduction in economic hours, for boys the amount of time spent on economic activities (which mainly constitutes of agricultural work) increases (Table 2).

Further, examination of the results by gender reveals that the one robustly significant effect of child labor for girls is on the probability of marriage. There are several possible interpretations of this. For instance, similarly to boys, it is possible that the type of incremental work induced by our instrument (chores for girls) makes girls more valuable on the marriage market.

VI. Extensions and Robustness Checks

VI.1 Threshold effects in child labor hours

In this section we consider whether the negative effects of child labor estimated in Table 6 are truly linear or whether there are any threshold effects, i.e., levels of child labor beyond which we begin to see truly harmful effects but below which child labor is not particularly harmful. Table 7 presents five alternative specifications in which the child labor variable is coded as an indicator for having worked more than a specified cutoff: 0 hours per week, 6.5 hours per week (the 25th percentile), 15 hours per week (the median), 24 hours per week (the 75th percentile), and 35 hours per week (the 90th percentile).

It is noteworthy that simply working any level of positive hours does not have a statistically significant impact on the majority of the outcomes, other than marriage and labor productivity. However, even in the instance of these two outcomes the effect of any level of positive hours is only marginally significant at a 10 percent level. Of course, it must be noted that less than 10 percent of the population do not work, so the effect may not be estimated precisely. When child labor is coded as working more than 6.5 hours, its impact on marriage probability and labor productivity becomes more robust, with effects that are significant at the 1 percent level. The negative effect of child labor on education outcomes, found in the previous specification (Table 6) also begins to come through at this level of child labor, though it only becomes significant at above median hours (over 15 per week) of work. At this threshold, working in childhood reduces time spent in schooling by 2.6 years, and reduces the probability of completing primary school by 36 percent. Increasing the threshold for the child labor indicator further increases the magnitude of the negative effect on completed schooling from 2.6 to 3.3 to 4.5 years. This exercise suggests that there are significant, negative effects of child labor from

even moderate levels of intensity, but that the effects increase in magnitude with the intensity of child labor.

VI.2 Robustness checks: the full sample of children and adults

In this section, we examine the sensitivity of our results to the choice of sample. As discussed in Section IV.1, the regressions in Tables 2-7 are estimated on the sample of 7 to 15 year olds who were either in school as of each wave or who were not in school but still young enough to start school at some later point (under the age of 10). This restriction is particularly important for education outcomes because, while child labor and education are simultaneous decisions, we are interested in identifying the impact of child labor (rather than, say, delayed enrollment) on educational attainment. Including children who are unlikely ever to attend school (i.e. those older than 9 and not at school in wave 4) in the sample would naturally lead us toward finding a stronger negative impact of child labor on education.

When we run our instrumented regression on the full sample of children between ages 7 and 15 (Table 8, row 1), we find, as expected, larger coefficients for school years and primary school completion. While the marriage result is unchanged, the negative effect on the marginal productivity of labor is 9 percentage points smaller. Given that we have added a number of children who are unlikely to have ever attended school for any substantial amount of time, the reduced negative effect on labor productivity suggests that the negative effect we were obtaining with the restricted sample could reflect displaced education.

Next, we turn to the adult sample (individuals aged 20 and older in wave 1) as a comparison group for the child sample. In particular, we estimate the same two-stage least squares specification used for children on the adult sample. In Table 8, row 2, we see that the

effect of labor instrumented using this two-stage least squares specification on outcomes of adults is not statically significant. This result suggests that the instruments do not affect the outcomes of interest through any channels other than child labor, which is consistent with the behavior of instruments that satisfy the exclusion restriction.

VI.3 Robustness: magnitude of the shock

Finally, we exploit cross-sectional variation in the size of shocks to test whether our second stage results are driven by large shocks or whether we obtain similar results for smaller shocks. As discussed in Section IV.2, we are more confident that the exclusion restriction required for a valid instrumental variable (namely that the instrument affects the outcome only through the endogenous variable) will be satisfied for small shocks, since these are less likely to have direct long lasting effects on household outcomes. In Table 9, we re-estimate our results, using three alternative definitions of the shock. We set the shock magnitude thresholds as shocks that result in a loss of 5, 10 and 20 percent of the crop and excluding those individuals who experienced shocks larger than the respective thresholds. For instance, at the 20 percent threshold we exclude 69 individuals from our original sample who experienced a crop loss greater than 20 percent. We find that the magnitude of the impact of child labor on the outcomes does not change substantially with magnitude of shock, suggesting that our findings are not driven by long-lasting effects of larger shocks and that the exclusion restriction is satisfied.

VII. Conclusion and Future Research

In this paper we investigated the impact of child labor on education and labor market outcomes using panel data from the Kagera region of Tanzania. Building on our previous work,

we used the occurrence of crop and rainfall shocks as instruments for child labor. In instrumental variables specifications, we find a negative and significant effect of child labor on school years, the probability of completing primary school, and marginal labor productivity 10 to 13 years later. Moreover, child labor is significantly positively associated with the probability of being married, which, in the context of this sample, is equivalent to an increased probability of earlier marriage.

The education results are mainly driven by the sample of boys. Working in childhood among boys is also found to increase the probability of being involved in farming 10 to 13 years later. For girls we find a robustly positive effect on the probability of marriage. In conjunction with our finding that the extra child labor induced by crop shocks is primarily chores, this suggests the possibility that child labor pushes girls away from agriculture and toward household activities and marriage. In future work we intend to use data on bride prices to examine whether child labor in fact increases girls' value on the marriage market.

References

- Akabayashi, H., and Psacharopoulos, George. (1999) "The Trade-off between Child Labor and Human Capital: A Tanzanian Case." *Journal of Development Studies* 35(5): 120-140.
- Basu, Kaushik (1999) "Child Labor: Cause, Consequence, and Cure with Remarks on International Labor Standards." *Journal of Economic Literature* 37: 1083-1119.
- Beegle, Kathleen, Rajeev Dehejia, and Roberta Gatti (2006) "Child Labor and Agricultural Shocks." *Journal of Development Economics* 81(1): 80-96.
- Beegle, Kathleen, Rajeev Dehejia, and Roberta Gatti (2005) "Why Should We Care about Child Labor? The Education, Labor Market, and Health Consequences of Child Labor." World Bank Policy Research Working Paper 3479. CEPR Discussion Paper 4443. NBER Working Paper No. 10980.
- Beegle, Kathleen, Joachim De Weerd, and Stefan Dercon (2006a) "Kagera Health and Development Survey 2004 Basic Information Document." World Bank.
- Beegle, Kathleen, Joachim De Weerd, Stefan Dercon (2006b) "Poverty and Wealth Dynamics in Tanzania: Evidence from a Tracking Survey." manuscript.
- Behrman, Jere R., Alexis Murphy, Agnes Quisumbing, Usha Ramakrishna, and Kathryn Young. (2006) "What is the Real Impact of Schooling on Age of First Union and Age of First Parenting? New Evidence from Guatemala." World Bank Policy Research Working Paper no. 4023.
- Bezerra, Márcio Eduardo G., Ana Lúcia Kassouf, Mary Arends-Kuenning. 2007. "The Impact of Child Labor and School Quality on Academic Achievement in Brazil." mimeo.
- Boozer, Michael, and T. Suri (2001) "Child Labor and Schooling Decisions in Ghana." Working paper, Yale University.
- Cavalieri, C. (2002) "The Impact of Child Labor on Educational Performance: An Evaluation of Brazil." manuscript.
- Emerson, P., and A. Portela Souza (2007) "Is Child Labor Harmful? The Impact of Working Earlier in Life on Adult Earnings." IZA Discussion Paper No. 3027.
- Heady, C. (2003) "The Effect of Child Labor on Learning Achievement." *World Development* 31: 385-398.
- Horowitz, Andrew, and Jian Wang (2004) "Favorite Son? Specialized Child Laborers and Students in Poor LDC Households." *Journal of Development Economics* 73: 631-642.
- International Labour Organization (2003) *Investing in Every Child: An Economic Study of the Costs and Benefits of Eliminating Child Labour*. Geneva: International Labour Office.

International Labour Organization (2002). *A Future Without Child Labour*. Geneva: International Labour Office.

Jacoby, H. (1993). “Shadow Wages and Peasant Family Labor Supply: An Econometric Application to the Peruvian Sierra,” *Review of Economic Studies* 60(4): 903-921.

Krutikova, Sofya (2006) “Impact of Child Labor on Education Attainment in Adulthood: Evidence from Rural Tanzania.” manuscript.

Morduch, Jonathan (1995) “Income Smoothing and Consumption Smoothing.” *Journal of Economic Perspectives* 9 (3): 103-14.

O’Donnell O., Furio Rosati, and E. Van Doorsaler (2005) “Health Effects of Child Work: Evidence from Rural Vietnam.” *Journal of Population Economics* 18(3): 437-467.

Patrinos, Harry A., and George Psacharopoulos (1995) “Educational Performance and Child Labor in Paraguay.” *International Journal of Educational Development* 15: 47–60.

Patrinos, Harry A., and George Psacharopoulos (1997) “Family Size, Schooling and Child Labor in Peru—An Empirical Analysis.” *Journal of Population Economics* 10: 387– 405.

Pörtner, Klaus (2006) “Gone with the Wind? Hurricane Risk, Fertility and Education.” Manuscript.

Ravallion, Martin and Michael Lokshin (2005) “Lasting Local Impacts of an Economywide Crisis.” World Bank Policy Research Working Paper No. 3506.

Ravallion, Martin and Quentin Wodon (2000) “Does Child Labour Displace Schooling? Evidence on Behavioral Responses to an Enrollment Subsidy.” *The Economic Journal* 110: 158-175.

Ray, R., and G. Lancaster (2004) “The Impact of Children’s Work on Schooling: Multi Country Evidence Based on SIMPOC Data.” manuscript.

Rosenzweig, Mark and Ken Wolpin (1985) “Specific Experience, Household Structure and Intergenerational Transfers: Farm Family Land and Labor Arrangements in Developing Countries.” *Quarterly Journal of Economics* 100: 961-987.

World Bank (2004) “User’s Guide to the Kagera Health and Development Survey Datasets.” mimeo.

World Bank. (2007) “World Development Report 2007: Development and the Next Generation.” World Bank.

Table 1: Summary Statistics

	Mean	SD
<u>Baseline sample</u>		
Hours worked in last 7 days	16.79	13.42
Chore hours in last 7 days	10.54	9.05
Crop shock	0.34	-
Female	0.49	-
Age	10.91	2.60
Number of observations		4,746
<u>Panel sample: 1991-2004</u>		
<i>Baseline</i>		
Mean hours	16.79	10.55
Mean hours (predicted)	16.57	5.73
Female	0.49	-
Age (wave 4)	11.58	2.98
Mother's education 1-6 years	0.36	-
Mother's education 7 years	0.32	-
Mother's education 8+ years	0.02	-
Father's education 1-6 years	0.43	-
Father's education 7 years	0.32	-
Father's education 8+ years	0.13	-
Expenditure per capita (log Tsh) (wave 1)	10.92	0.81
Asset value per capita (log Tsh) (wave 1)	10.05	1.60
Land value per capita (log Tsh) (wave 1)	10.03	2.26
Any crop shock (waves 1-4)	0.69	-
Rainfall deviation	0.21	0.54
<i>2004</i>		
School years	6.36	2.77
Completed primary	0.78	-
Stayed in/near village	0.69	-
Farming in past 12 months	0.76	-
Growing cash crop	0.55	-
Wage/salary job in past 12 months	0.37	-
Married	0.51	-
Marginal product of labor (log Tsh/day)	7.03	0.87
Number of observations		1,311

Notes: Baseline sample is restricted to children in school at baseline or less than 10 years of age and not yet enrolled. It includes children who are measured up to 4 times in the baseline panel (1991-1994). Hours includes hours working in economic (income generating) activities and in chores. Panel sample is the subset of children in baseline sample who are re-interviewed in 2004.

Table 2: 1st Stage Estimation of Child Labor Hours

	(1)	(2)	(3) Girls age 7-10	(4) Girls age 11-15	(5) Boys age 7-10	(6) Boys age 11-15
Crop shock	0.487 (0.538)	1.221 (2.141)	-1.534* (0.853)	2.198** (0.961)	-0.110 (0.858)	0.800 (0.839)
Crop shock * Age		-0.072 (0.192)				
Crop shock * Female		-14.608*** (2.589)				
Crop shock * Female * Age		1.350*** (0.245)				
Rainfall deviation	0.901 (0.596)	0.951* (0.589)	0.372 (0.907)	1.984* (1.098)	0.591 (0.990)	0.875 (0.994)
Age	6.292*** (0.715)	5.918*** (0.727)	-7.750 (5.406)	4.706 (5.805)	5.898 (6.446)	8.277 (5.267)
Age squared	-0.208*** (0.033)	-0.201*** (0.033)	0.629** (0.319)	-0.124 (0.222)	-0.176 (0.379)	-0.306 (0.202)
Female	1.852*** (0.422)	1.830*** (0.500)				
Number of observations	4,746	4,746	1,066	1,259	1,065	1,356

Notes: Regressions from waves 1-4 at baseline for restricted sample of children described in text ages 7-15. Additional controls include age, age squared, household size, mother's education, father's education, log value of per capita household land, asset and expenditure, season dummies, and region dummies. Standard errors are in parentheses. *** indicates significance at 1%; ** at 5%; and, * at 10%. Work hours include hours working in economic (income generating) activities and in chores.

Table 3: Correlation between Instruments and Lagged Instruments and Household Characteristics

	(1)	(2)	(3)
	Crop shock	Crop shock	Rainfall deviation
Lagged shock, t-1		0.101 (0.075)	
Lagged shock t-2		0.043 (0.034)	
Lagged shock t-3		0.029 (0.042)	
Household size	-0.000 (0.003)	0.003 (0.005)	0.002 (0.005)
Father's education 1-6 years	-0.003 (0.029)	0.000 (0.038)	-0.021 (0.041)
Father's education 7 years	-0.030 (0.030)	0.016 (0.047)	-0.008 (0.042)
Father's education 8+ years	0.042 (0.036)	-0.018 (0.044)	-0.016 (0.045)
Mother's education 1-6 years	0.019 (0.023)	-0.018 (0.038)	0.001 (0.036)
Mother's education 7 years	-0.029 (0.023)	-0.040 (0.041)	0.005 (0.037)
Mother's education 8+ years	-0.013 (0.076)	-0.056 (0.055)	0.216* (0.118)
Number of observations	4,746	651	4,746

Notes: Regressions from waves 1-4 at baseline for restricted sample of children described in text ages 7-15. Additional controls include age, age squared, female, and season dummies (or the village means in the case of column 2). *** indicates significance at 1%; ** at 5%; and, * at 10%.

Table 4: Long-run Shock Effect on Household Wealth in 2004

	(1)	(2)	(3)	(4)	(5)
	Physical assets	Business assets	Durables assets	Farm equipment	Land
Crop shock in waves 1-4	3325 (2884)	312 (347)	26 (64)	16 (15)	2829 (2842)
Rainfall deviation (mean waves 1-4)	-4797 (3913)	-479 (407)	-95 (58)	-39 (26)	-3455 (3897)
Number of observations	1,263	1,263	1,263	1,263	1,263

Notes: Dependent variables measured in 1,000 Tshilling values in 2004. Controls include age, age squared, female, mother's education, father's education, household size, and region dummies. Standard errors are clustered at the wave 1-4 household level.

Table 5: Impact of Child Labor in Waves 1-4 on Outcomes in Wave 5: OLS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School years	Completed primary	Stayed in/near village	Farming in past 12 months	Growing cash crop	Wage/salary job in past 12 months	Married	Marginal productivity of labor
<i>Boys and girls</i>								
Mean child labor hours, waves 1-4	-0.008 (0.009)	-0.001 (0.001)	-0.001 (0.001)	0.001 (0.001)	0.002 (0.001)	-0.001 (0.001)	0.002 (0.001)	0.002 (0.003)
% impact on outcome of 1 sd ↑ in hours	-0.7%	-1.3%	-1.4%	1.3%	1.8%	-2.7%	2.0%	0.1%
Level impact on outcome of 1 sd ↑ in hours	-0.05	-0.01	-0.01	0.01	0.01	-0.01	0.01	0.01
Number of observations	1,311	1,311	1,311	1,311	1,311	1,311	1,311	850
<i>Girls</i>								
Mean child labor hours, waves 1-4	-0.010 (0.014)	-0.002 (0.002)	-0.003 (0.002)	0.0001 (0.002)	0.001 (0.002)	0.0001 (0.002)	0.004** (0.002)	0.001 (0.004)
% impact on outcome of 1 sd ↑ in hours	-0.9%	-1.3%	-2.6%	0.08%	1.0%	0.2%	4.5%	0.09%
Level impact on outcome of 1 sd ↑ in hours	-0.06	-0.01	-0.02	0.00	0.00	0.00	0.02	0.00
Number of observations	639	639	639	639	639	639	639	399
<i>Boys</i>								
Mean child labor hours, waves 1-4	-0.007 (0.012)	-0.0001 (0.002)	0.001 (0.002)	0.002 (0.002)	0.003 (0.002)	-0.002 (0.002)	-0.001 (0.002)	0.007* (0.004)
% impact on outcome of 1 sd ↑ in hours	-0.6%	-0.07%	0.8%	1.6%	3.2%	-2.2%	-1.6%	0.5%
Level impact on outcome of 1 sd ↑ in hours	-0.04	-0.00	0.00	0.01	0.02	-0.01	-0.01	0.04
Number of observations	672	672	672	672	672	672	672	451

Notes: Each cell presents results from a separate regression, with a common specification across sets of rows: all boys and girls between 7 and 15 who are in school or can enroll; all girls satisfying the same age and school criteria; and all boys satisfying the same age and school criteria. The magnitude of the effect is computed as the percentage and level impact on the outcome of a one standard deviation (5.7 hour) increase in child labor. Additional controls include a sex dummy, age, age squared, mother's education, father's education, region dummies, wave 1 log expenditure, wave 1 log assets, and wave 1 log land holdings.

Table 6: Impact of Child Labor: 2SLS of 2004 outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School years	Completed primary	Stayed in/near village	Farming in past 12 months	Growing cash crop	Wage/salary job in past 12 months	Married	Marginal productivity of labor
<u>Boys and girls</u>								
Mean child labor hours, waves 1-4	-0.076** (0.033)	-0.012** (0.005)	-0.008 (0.005)	0.007 (0.005)	-0.003 (0.006)	0.001 (0.005)	0.014*** (0.005)	-0.036** (0.014)
% impact on outcome of 1 sd ↑ in hours	-6.8%	-8.8%	-6.6%	5.3%	-3.1%	1.5%	15.7%	-2.9%
Level impact on outcome of 1 sd ↑ in hours	-0.44	-0.07	-0.05	0.04	-0.02	0.01	0.08	-0.21
Number of observations	1,311	1,311	1,311	1,311	1,311	1,311	1,311	850
<u>Girls</u>								
Mean child labor hours, waves 1-4	-0.034 (0.042)	-0.006 (0.006)	-0.006 (0.007)	-0.007 (0.007)	-0.007 (0.008)	0.006 (0.007)	0.017*** (0.007)	-0.024 (0.016)
% impact on outcome of 1 sd ↑ in hours	-3.1%	-4.4%	-5.0%	-5.1%	-7.0%	17.2%	14.5%	-2.1%
Level impact on outcome of 1 sd ↑ in hours	-0.19	-0.03	-0.03	-0.04	-0.04	0.03	0.10	-0.14
Number of observations	639	639	639	639	639	639	639	399
<u>Boys</u>								
Mean child labor hours, waves 1-4	-0.143** (0.057)	-0.019** (0.009)	-0.005 (0.008)	0.023*** (0.008)	0.006 (0.009)	-0.003 (0.009)	0.020*** (0.008)	-0.035 (0.022)
% impact on outcome of 1 sd ↑ in hours	-12.9%	-14.1%	-3.9%	18.3%	6.4%	-3.3%	31.8%	-2.7%
Level impact on outcome of 1 sd ↑ in hours	-0.82	-0.11	-0.03	0.13	0.03	-0.02	0.11	-0.20
Number of observations	672	672	672	672	672	672	672	451

Notes: Each cell presents results from a separate regression, with a common specification across sets of rows: all boys and girls between 7 and 15 who are in school or can enroll; all girls satisfying the same age and school criteria; and all boys satisfying the same age and school criteria. The magnitude of the effect is computed as the percentage and level impact on the outcome of a one standard deviation (5.7 hour) increase in child labor. Additional controls include a sex dummy, age, age squared, mother's education, father's education, region dummies, wave 1 log expenditure, wave 1 log assets, and wave 1 log land holdings. Block boot-strapped standard errors are in parentheses.

Table 7: 2SLS robustness checks: threshold effects of child labor hours

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School years	Completed primary	Stayed in/near village	Farming in past 12 months	Growing cash crop	Wage/salary job in past 12 months	Married	Marginal productivity of labor
<u>Works more than 0 hours</u>								
Mean child labor hours, waves 1-4	1.187	0.070	-0.660	0.096	0.619	-0.319	0.984*	-2.382*
	(3.560)	(0.551)	(0.564)	(0.538)	(0.603)	(0.530)	(0.548)	(1.282)
% impact on outcome of working>0 hours	18.7%	9.0%	-95.7%	12.6%	112.5%	-86.2%	192.9%	-33.9%
Level impact on outcome of working>0 hours	1.19	0.07	-0.66	0.10	0.62	-0.32	0.98	-2.38
Number of observations	1,311	1,311	1,311	1,311	1,311	1,311	1,311	850
<u>Works more than 6.5 hours (25th percentile)</u>								
Mean child labor hours, waves 1-4	-1.743	-0.315	-0.545**	0.146	-0.080	0.098	0.775***	-1.922***
	(1.725)	(0.248)	(0.274)	(0.258)	(0.296)	(0.301)	(0.270)	(0.659)
% impact on outcome of working>6.5 hours	-27.4%	-40.4%	-79.0%	19.2%	-14.5%	26.5%	152.0%	-27.3%
Level impact on outcome of working>6.5 hours	-1.74	-0.32	-0.55	0.15	-0.08	0.10	0.78	-1.92
Number of observations	1,311	1,311	1,311	1,311	1,311	1,311	1,311	850
<u>Works more than 15 (median)</u>								
Mean child labor hours, waves 1-4	-2.554**	-0.356**	-0.225	0.263	-0.062	0.013	0.589***	-1.176***
	(1.099)	(0.159)	(0.163)	(0.170)	(0.202)	(0.183)	(0.171)	(0.428)
% impact on outcome of working>15 hours	-40.2%	-45.6%	-32.6%	34.6%	-11.3%	3.5%	115.5%	-16.7%
Level impact on outcome of working>15 hours	-2.55	-0.36	-0.23	0.26	-0.062	0.013	0.59	-1.18
Number of observations	1,311	1,311	1,311	1,311	1,311	1,311	1,311	850

Table 7 (continued): 2SLS robustness checks: threshold effects of child labor hours

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School years	Completed primary	Stayed in/near village	Farming in past 12 months	Growing cash crop	Wage/salary job in past 12 months	Married	Marginal productivity of labor
<i>Works more than 24 hours (75th percentile)</i>								
Mean child labor hours, waves 1-4	-3.298*** (1.233)	-0.497*** (0.184)	-0.220 (0.192)	0.293 (0.215)	-0.154 (0.237)	0.107 (0.206)	0.535*** (0.198)	-1.278** (0.503)
% impact on outcome of working>24 hours	-51.9%	-63.7%	-31.9%	38.6%	-28.0%	28.9%	105.0%	-18.2%
Level impact on outcome of working>24 hours	-3.3	-0.5	-0.22	0.29	-0.15	0.11	0.54	-1.23
Number of observations	1,311	1,311	1,311	1,311	1,311	1,311	1,311	850
<i>Works more than 35 hours (90th percentile)</i>								
Mean child labor hours, waves 1-4	-4.454** (1.935)	-0.776*** (0.268)	-0.355 (0.336)	0.412 (0.305)	-0.109 (0.356)	0.029 (0.294)	0.581** (0.284)	-1.609* (0.843)
% impact on outcome of working>35 hours	-70.0%	-99.5%	-51.4%	54.2%	-19.8%	7.8%	113.9%	-22.9%
Level impact on outcome of working>35 hours	-4.45	-0.78	-0.36	0.41	-0.11	-0.03	0.58	-1.61
Number of observations	1,311	1,311	1,311	1,311	1,311	1,311	1,311	850

Notes: Each cell presents results from a separate regression, with a common specification across sets of rows. The first set of rows uses an indicator for having worked more than 0 hours to measure child labor. The second set of rows uses an indicator for having worked 10 or more hours per week, and the third set of rows uses an indicator for having worked 25 or more hours per week. The magnitude of the effect is computed as the percentage and level impact on the outcome of a one standard deviation (5.7 hour) increase in child labor. Additional controls include a sex dummy, age, age squared, mother's education, father's education, region dummies, wave 1 log expenditure, wave 1 log assets, and wave 1 log land holdings. Block boot-strapped standard errors are in parentheses.

Table 8: 2SLS robustness checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School years	Completed primary	Stayed in/near village	Farming in past 12 months	Growing cash crop	Wage/salary job in past 12 months	Married	Marginal productivity of labor
<u>All kids sample</u>								
Mean child labor hours, waves 1-4	-0.097*** (0.036)	-0.015** (0.006)	-0.007 (0.005)	0.005 (0.005)	-0.003 (0.005)	0.004 (0.004)	0.012** (0.005)	-0.027*** (0.009)
% impact on outcome of 1 sd ↑ in hours	-9.3%	-11.8%	-5.8%	3.8%	-3.1%	6.2%	12.7%	-2.2%
Level impact on outcome of 1 sd ↑ in hours	-0.556	-0.086	-0.04	0.029	-0.017	0.023	0.069	-0.155
Number of observations	1,460	1,460	1,460	1,460	1,460	1,460	1,460	952
<u>Adults (20 +)</u>								
Mean child labor hours, waves 1-4	-0.039 (0.033)	-0.006 (0.004)	-0.003 (0.004)	-0.001 (0.004)	-0.001 (0.005)	0.005 (0.004)	0.002 (0.002)	-0.001 (0.008)
% impact on outcome of 1 sd ↑ in hours	-4.9%	-7.2%	-2.0%	-0.7%	-0.8%	10.2%	1.2%	-0.1%
Level impact on outcome of 1 sd ↑ in hours	-0.223	-0.034	-0.017	-0.006	-0.006	0.029	0.011	-0.006
Number of observations	1,525	1,525	1,525	1,525	1,525	1,525	1,525	1,126

Notes: Each cell presents results from a separate regression, with a common specification across sets of rows. The all kids sample uses all children between 7 and 15 (hence adds those children 7 to 15 not in school in waves 1 to 4 excluded from the main specification). The adult sample uses all adults age 20 or more in waves 1 to 4. The final row presents the difference between the two estimates and the standard error of the difference. The magnitude of the effect is computed as the percentage and level impact on the outcome of a one standard deviation (5.7 hour) increase in child labor. Additional controls include a sex dummy, age, age squared, mother's education, father's education, region dummies, wave 1 log expenditure, wave 1 log assets, and wave 1 log land holdings. Block boot-strapped standard errors are in parentheses.

Table 9: 2SLS robustness checks: magnitude of the shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	School years	Completed primary	Stayed in/near village	Farming in past 12 months	Growing cash crop	Wage/salary job in past 12 months	Married	Marginal productivity of labor
<u>Crop shocks < 5% of total</u>								
Mean child labor hours, waves 1-4	-0.089** (0.042)	-0.012** (0.006)	0.002 (0.006)	0.013** (0.006)	-0.001 (0.007)	-0.004 (0.006)	0.023*** (0.007)	-0.028* (0.017)
% impact on outcome of 1 sd ↑ in hours	-8.0%	-8.8%	1.7%	9.8%	-1.0%	-6.2%	25.8%	-2.3%
Level impact on outcome of 1 sd ↑ in hours	-0.51	-0.069	0.011	0.074	-0.006	-0.023	0.132	-0.16
Number of observations	1,189	1,189	1,189	1,189	1,189	1,189	1,189	776
<u>Crop shocks < 10% of total</u>								
Mean child labor hours, waves 1-4	-0.089** (0.041)	-0.012** (0.006)	-0.001 (0.006)	0.011* (0.006)	-0.001 (0.006)	-0.004 (0.006)	0.021*** (0.007)	-0.027** (0.013)
% impact on outcome of 1 sd ↑ in hours	-8.0%	-8.8%	-0.8%	8.3%	-1.0%	-6.2%	23.6%	-2.2%
Level impact on outcome of 1 sd ↑ in hours	-0.51	-0.069	-0.006	0.063	-0.006	-0.023	0.12	-0.155
Number of observations	1,214	1,214	1,214	1,214	1,214	1,214	1,214	791
<u>Crop shocks < 20% of total</u>								
Mean child labor hours, waves 1-4	-0.092** (0.039)	-0.013** (0.006)	-0.001 (0.006)	0.009 (0.006)	-0.002 (0.007)	-0.000 (0.006)	0.021*** (0.008)	-0.039** (0.017)
% impact on outcome of 1 sd ↑ in hours	-8.3%	-9.6%	-0.8%	6.8%	-2.1%	-	23.6%	-3.2%
Level impact on outcome of 1 sd ↑ in hours	-0.53	-0.074	-0.006	0.052	-0.011	-	0.12	-0.223
Number of observations	1,242	1,242	1,242	1,242	1,242	1,242	1,242	809

Notes: Each cell presents results from a separate regression, with a common specification across sets of rows. The second set of rows codes the instrument as 0 if the household did not experience a shock and 1 if it experienced a shock of 5 per cent or less of total crop. The second set of rows codes the instrument as 0 if the household did not experience a shock and 1 if it experienced a shock of 10 per cent or less of total crop. The second set of rows codes the instrument as 0 if the household did not experience a shock and 1 if it experienced a shock of 20 per cent or less of total crop. Each set of rows excludes those observations with shocks greater than the specified cutoff. The magnitude of the effect is computed as the percentage and level impact on the outcome of a one standard deviation (5.7 hour) increase in child labor. Additional controls include a sex dummy, age, age squared, mother's education, father's education, region dummies, wave 1 log expenditure, wave 1 log assets, and wave 1 log land holdings. Block boot-strapped standard errors are in parentheses.