ELSEVIER

Journal of Development Economics 81 (2006) 80–96

JOURNAL OF Development ECONOMICS

www.elsevier.com/locate/econbase

Child labor and agricultural shocks☆

Kathleen Beegle [a], Rajeev H. Dehejia [b,c,*], Roberta Gatti [a,d]

[a] *Development Research Group, The World Bank, United States*
[b] *Department of Economics and SIPA, Columbia University, United States*
[c] *NBER, United States*
[d] *CEPR, United States*

Received 15 October 2004; received in revised form 3 May 2005; accepted 27 May 2005

Abstract

This paper examines the relationship between household income shocks and child labor. In particular, we investigate the extent to which transitory income shocks lead to increases in child labor and whether household asset holdings mitigate the effects of these shocks. Using data from a household panel survey in Tanzania, we find that both relationships are significant. We investigate mechanisms that could account for these results, including buffer stocks and borrowing.

JEL classification: J82; J22; G20; O16
Keywords: Child labor; Buffer stocks; Credit constraints; Consumption smoothing

1. Introduction

This paper examines the relationship between household income shocks and child labor. In particular, we investigate the extent to which transitory income shocks lead to increases in child labor and whether household asset holdings mitigate the effects of these shocks. Using data from a household panel survey in Tanzania, we find that both

☆ The views expressed here do not necessarily reflect those of the World Bank or its member countries.

* Corresponding author. Department of Economics and SIPA, Columbia University, United States. Tel.: +1 212 854 4659.
E-mail address: rd247@columbia.edu (R.H. Dehejia).

0304-3878/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.jdeveco.2005.05.003

of these factors are significantly related to the use of child labor. We investigate mechanisms that could account for these results, including buffer stocks and borrowing.

The question and our results are important for a number of reasons. First, they point to a significant determinant of child labor and to mechanisms that can potentially be used to reduce it. Traditionally, child labor has been viewed as a consequence of poverty (see for example Krueger, 1996, and Fallon and Tzannatos, 1998). This relationship has been put into question by a number of recent within-country studies (see inter alia Canagarajah and Nielsen, 1999; Grootaert and Patrinos, 1999; Boozer and Suri, 2001; Bhalotra and Heady, 2003). Our work examines the role that child labor plays as a buffer against transitory shocks, and suggests that insurance or access to credit might reduce the extent of child labor.

Second, our work explores the role of asset holdings in determining a household's response to shocks, and as such relates to both the recent theoretical literature on the role of credit constraints in explaining child labor (Baland and Robinson, 2000; Ranjan, 2001) and the buffer stock literature (Deaton, 1992).[1] From the household's point of view, child labor entails a trade-off between immediate benefits (increased current income) and, to the extent it interferes with the accumulation of the child's human capital, potential long-run costs (lower future earnings potential).[2] Thus, when faced with a transitory shock, households would use asset holdings (either as a buffer or as collateral against credit) to offset the shock.

Third, our work relates to the literature on the permanent income hypothesis and consumption smoothing (see, inter alia, Zeldes, 1989; Townsend, 1994; Chaudhuri and Ravallion, 1997; Morduch, 1994, 1995, 1999). If households succeed in smoothing their consumption profile, but lack buffer stocks or are credit constrained, they are forced to resort to other mechanisms to cope with income shocks. This paper examines one such mechanism, namely child labor. Related work has examined the response of schooling (Jacoby, 1994; Jacoby and Skoufias, 1997; Fitzsimmons, 2003) and of parental labor supply (Kochar, 1999) to income shocks.[3]

Using four rounds of household panel data from the Kagera region of Tanzania, we show that transitory income shocks — as measured by accidental crop loss — lead to significantly increased child labor. In the Kagera region, households use almost no purchased inputs and very rudimentary technology (only the most basic tools; no tractors for example), and the use of wage labor is very limited. When hit by a shock, we show that households tend to increase their use of child labor, typically by having children

[1] Credit market imperfections can also impact child health. For example, Foster (1995) finds that credit market imperfections influence growth patterns for children in landless households in Bangladesh.

[2] In outlining child labor issues and directions for the World Bank, Fallon and Tzannatos (1998) point out that child labor can have additional costs in terms of harmful effects on physical health and mental well-being (such as psychological and social adjustment) of children. Moreover, others contend that the working conditions for children are far below those of adults in terms of hours worked, wages, and safety.

[3] It is not obvious that for children there is a one-for-one trade-off between time spent in school and time spent working. Hours in either activity may be sufficiently low on average that an increase in time spent in one activity will simply crowd out leisure time (see Ravallion and Wodon, 2000).

substitute adult labor in household activities such as gathering firewood and water. In parallel, we also show a decrease in school attendance. For all of these effects, we find that households with assets are better able to offset shocks. We find that the level of household asset holdings decreases in response to shocks — consistent with a buffer stock interpretation of our results — but that wealthier households draw down their assets to a lesser extent, suggesting that these households could be borrowing in response to shocks. We present evidence on household loans consistent with this view, but also discuss alternative interpretations of our results.

Despite its theoretical centrality, this is among the first studies explicitly examining this relationship. Using cross-country data, Dehejia and Gatti (in press) find a significant negative relationship between child labor and access to credit (as proxied by financial development). Guarcello et al. (2003), using data from Guatemala, find that child labor increases in response to broadly defined income shocks and self-reported credit rationing. Their results, though suggestive, do not control for other household level characteristics that could be correlated with credit (e.g., household fixed effects), and do not examine the interaction between shocks and credit. Edmonds (2002) uses an identification strategy that is related to, though different from, our own. Whereas our identification is based on finding an effect of an unanticipated transitory income shock, Edmonds examines an anticipated permanent increase in income (an increase in old-age pension support in South Africa). The two approaches are clearly complementary. A caveat to Edmonds' approach is that households may fail to borrow against their future old-age pension simply because they are myopic. In this case, the effect he estimates would simply be a wealth effect.

The paper is organized as follows. The empirical strategy is outlined in Section 2. In Section 3, we describe our data. In Section 4, we present our results, and Section 5 concludes.

2. Empirical strategy and specification

We discuss first the specifications we use to examine the effect of income shocks on child labor, and then return to the question of whether or not the income shocks we examine are plausibly transitory.

We examine the effect of shocks on child labor hours. Our basic specification is:

$$y_{ijt} = \beta_0 + \beta_1 X_{ijt} + \beta_2 shock_{ijt} + \varepsilon_{ijt} \quad (1)$$

where: subscripts index individuals (i), households (j), and survey rounds ($t = 1,\ldots,T$); y is child labor hours (as well as education and specific child labor activities); *shock* is our measure of income shocks (an indicator of crop lost to pests and other calamities, discussed in detail in the next section); and X contains a set of controls including individual, household, and community characteristics. We anticipate transitory shocks to lead to an increase in child labor if insurance is incomplete, i.e. we expect $\beta_2 > 0$ for child labor (and $\beta_2 < 0$ for schooling).

To investigate the role of assets in buffering against shocks, we examine whether the effect of shocks varies with households' asset holdings. In particular, we estimate the following specification:

$$y_{ijt} = \beta_0 + \beta_1 X_{ijt} + \beta_2 shock_{ijt} + \beta_3 \left(shock_{ijt} \cdot assets_{ijt}\right) + \beta_4 assets_{ijt} + \varepsilon_{ijt} \quad (2)$$

where *assets* measure non-land household wealth. The effect of interest is β_3, which captures the differential impact of a shock among households with different levels of assets. We expect asset holdings to mitigate the effect of shocks, i.e. $\beta_3 < 0$ for child labor (and $\beta_3 > 0$ for schooling).

Our measure of assets could be correlated with omitted variables that in turn are associated with reduced child labor, for example the labor intensity of a household's production technology or the extent of its social networks. Examples such as these motivate our inclusion of household fixed effects in the above specifications:

$$y_{ijt} = \alpha_j + \delta_t + \gamma_w + \beta_1 X_{ijt} + \beta_2 shock_{ijt} + \varepsilon_{ijt} \quad (3)$$

$$y_{ijt} = \alpha_j + \delta_t + \gamma_w + \beta_1 X_{ijt} + \beta_2 shock_{ijt} + \beta_3 \left(shock_{ijt} \cdot assets_{ijt}\right) + \beta_4 assets_{ijt} + \varepsilon_{ijt} \quad (4)$$

where α_j, δ_t, γ_w, are household, time (season), and survey round fixed effects respectively. Including household fixed effects also implies that crop shocks are measured relative to a household level (thus, by construction are mean zero). Another implication is that we are examining the effect of idiosyncratic risk, after aggregate (i.e., community) and non-time varying household risk factors have been purged.[4]

Finally, we return to the question of whether the shocks that we consider are plausibly exogenous and transitory. With respect to the first question, the concern is that the households who are most affected by shocks may systematically differ from those that are not. For example, Morduch (1994) discusses the possibility that poorer households may be more likely to experience shocks. We address this concern by controlling throughout for a wide range of time varying household characteristics and household fixed effects. Controls include mother's and father's schooling, the presence of the mother and father in the household, the child's age, and interview-wave and season fixed effects. We also consider whether households with high levels of child labor are more likely to experience shocks; in particular, we examine whether lagged child labor predicts which households experience shocks:

$$Pr\left(shock_{ijt} = 1\right) = f\left(\ childlabor_{ijt-1}, X_{jt}\right). \quad (5)$$

Regarding the transitory nature of the shocks, the concern then is not whether (lagged) household characteristics predict crop shocks (since we control for these), but

[4] In comparing our framework to Jacoby and Skoufias (1997) or Fitzsimmons (2003), it is worth underlining this point. Their measures of aggregate and idiosyncratic risk are fully absorbed by our household fixed effects.

whether households that have experienced a shock are more likely to experience shocks in the future. If this were the case, then a household experiencing a shock would anticipate a permanent reduction in income. We investigate this empirically by examining:

$$Pr(shock_{ijt} = 1) = f(shock_{ijt-1}, X_{jt}). \quad (6)$$

If households that have experienced shocks in the past are not more likely to experience them in the future, it lends credence to our claim that crop shocks are transitory. Finally, note that since our outcome specifications allow for fixed effects, this controls for an additional source of selection not captured in (6) and ensures that the shocks are mean zero (since the inclusion of fixed effects implies that we are implicitly taking out the household mean of shocks).

3. Data description and summary statistics

The data for this study are from a panel data set in the Kagera region in Tanzania. 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. Tree-crops and cassava, a commonly grown crop, have fairly continuous cultivation over the year.

The Kagera Health and Development Survey (KHDS) was part of a research project conducted by the World Bank and the University of Dar es Salaam (see inter alia Burke and Beegle, 2004; Ainsworth et al., 2002; Beegle, 2005).[5] The KHDS surveyed over 800 households in the region up to four times from 1991 to 1994 with an average interval between surveys of six to seven months.[6] Households are drawn from 51 communities, mostly villages, in the six districts of Kagera.

This data set has several features that make it particularly appropriate for the proposed analysis. First, the detailed household survey has a wide array of individual and household characteristics, including information on time use of all household members aged seven and older. This includes time spent in the previous week working in household businesses (farm and non-farm), in non-household businesses for wages, and in household chores. The household survey also includes information on crop loss, as well as measures of physical and financial assets in each of the four interviews.

Two potential concerns arise in using this data set. The first is attrition, which is a problem common to all longitudinal data sets. Overall, more than 90% of all households

[5] Studies on Tanzania, using other data sources, include Al-Samarrai and Peasgood (1998), Mason and Khandker (1997), Bommier and Lambert (2000), Al-Samarrai and Reilly (2000), and Burke (1998).

[6] The explicit objectives of the KHDS were to measure the economic impact of fatal illness (primarily due to HIV/AIDS) in the region and to propose cost-effective strategies to help survivors. For more information about this project, see Ainsworth et al. (1992) and World Bank (1993).

Table 1
Summary statistics

	(1)	(2)	(3)	(4)
Sample restriction: land acres		1–25.5	1–25.5	1–25.5
Sample restriction: shock			With shock	Without shock
Hours worked:				
Mean	18.13 (14.83)	18.2 (14.8)	18.4 (15.7)	18.0 (14.3)
Proportion > 0	0.90	0.90	0.89	0.91
log value of crop loss:				
Mean	2.96 (4.23)	3.04 (4.28)	8.66 (1.83)	0
Proportion > 0	0.34	0.35	1.00	0
log per capita assets	10.3 (1.55)	10.3 (1.47)	10.2 (1.51)	10.4 (1.45)
Per capita durables (10000s)	0.91 (8.20)	0.94 (8.87)	0.83 (8.04)	1.00 (9.29)
Per capita cash (10000s)	0.53 (4.95)	0.053 (5.33)	0.83 (8.50)	0.37 (2.16)
Per capita physical assets (10000s)	24.17 (241.20)	19.0 (66.8)	18.2 (67.9)	19.50 (66.19)
Per capita land value (10000s)	96.12 (39.89)	10.06 (41.65)	8.81 (18.11)	10.74 (49.95)
Farm household=1 if yes, else 0	0.76 (0.43)	0.79 (0.41)	0.78 (0.41)	0.79 (0.40)
Household size	7.83 (3.66)	7.96 (3.69)	7.92 (3.34)	7.99 (3.86)
Father's schooling: 1–6 years=1 if yes, else 0	0.42 (0.49)	0.44 (0.50)	0.45 (0.50)	0.44 (0.50)
Father's schooling: 7 years=1 if yes, else 0	0.33 (0.47)	0.33 (0.47)	0.31 (0.46)	0.34 (0.47)
Father's schooling: 8+ years=1 if yes, else 0	0.12 (0.32)	0.097 (0.30)	0.12 (0.32)	0.085 (0.28)
Mother's schooling: 1–6 years=1 if yes, else 0	0.35 (0.48)	0.35 (0.48)	0.38 (0.49)	0.34 (0.47)
Mother's schooling: 7 years=1 if yes, else 0	0.31 (0.46)	0.30 (0.46)	0.29 (0.45)	0.31 (0.46)
Mother's schooling: 8+ years=1 if yes, else 0	0.018 (0.13)	0.009 (0.093)	0.011 (0.10)	0.008 (0.086)
Observations	5591	4709	1653	3056

Standard deviations are in parentheses.

remained in the sample for all interviews. If children or households were to exit (or enter) our sample in a way that was correlated with shocks or child labor, this could bias our results. Employment is explicitly listed as the reason for departure for only two of the children in the sample. More generally, when we regress the probability of exiting from the sample on lagged labor, shocks, or asset holdings we find no significant relationship. We likewise find that these variables are not significant predictors of entry into the sample.

A second concern is that the survey design included an over-sample of households considered to be at risk based on the following criteria: those with a sick adult and/or those households that had experienced an adult death in the past 2 years at the time of the household listing (about 6 months before the first survey round). To the extent that these latent characteristics are time invariant, household fixed effects will address the problem. However, it is possible that households who have experienced a death in the previous two years respond differently to shocks. This is difficult to deal with econometrically, and must be acknowledged as a limitation of the analysis.

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 for two reasons. First, the concept of child labor (by

ILO standards) is not restricted to only economic activities.[7] Second, in the largely rural sample of households in this study, it may be difficult to distinguish time in household chore activities and time spent preparing subsistence food crops. For our study, we focus on children 7–15 years old. The upper age range for child labor studies is typically 14 or 15 years, the age of completed primary schooling if enrolled on time (and the minimum age for legal employment).

Table 1 presents summary statistics of the sample in our study, broken down by the three samples on which the regressions are run. In addition, the last two columns show summary statistics for the main sample separated into children in households that experienced an income shock and children in households that did not experience an income shock. In the pooled data, children worked on average 18 h in the previous week. Mean hours as well as most other covariates have a similar distribution in households with and without a shock. More than 90% of children worked at least 1 hour in the previous week. About one-third of children reside in households that report some crop loss. There are no significant differences in assets between households that experienced a shock and those that did not. The typical household size is quite large, over 7 members on average. In part, this reflects the sample of households with children. Levels of parental education are extremely low. Few children have fathers who had attended school beyond the primary level (12% of children) and almost none had mothers with more than primary schooling (2%).

Turning to our measure of income shock, for the identification strategy to be credible, we ideally want an income shock that is of a sufficient magnitude to affect household time allocation, exogenous to child labor decisions, and transitory. The data include reports of the value of crop loss due to insects, rodents, and other calamities (such as fire) in each survey round.[8] We use a dummy for positive crop losses as an indicator of shocks. Our measure of income shock has several advantages. First, since agriculture is the main economic activity in the Kagera region, many households experience shocks and these shocks are extremely relevant with respect to household income. Second, these shocks are plausibly exogenous and transitory. We provide evidence for both of these claims below.

We have a range of measures of household asset holdings: durable assets include the value of durable goods, such as radios, bicycles, fans, lamps and pots; physical assets include land (which is also separately measured), business equipment, and livestock; and finally we also observe cash holdings. In our main results, we use all non-land holdings (including cash) as our measure of assets; we exclude land because it is often positively associated with child labor demand and could confound our results. We subsequently study the effect of shocks on each kind of asset holding. Evidence from field interviews

[7] It 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 (perhaps impossible) to appropriately isolate the portion of time spent working on the farm that qualifies under this very nuanced definition.

[8] The questions on the survey are "Je, umepoteza sehemu ya mavuno yako kwa wadudu, panya, moto au kuharibika?" [Have you lost any part of the harvest to insects, rodents, fire, rotting or stolen?] and "Kamam ungeuza kiasi hicho wakati kilichoharibika ungepata fedha kiasi gani? (T-shillings)" [If you had sold this same quantity at the time you lost it, what is the most amount you could have gotten?].

suggests that durable assets are used as collateral for loans, while cash is more likely to be used as an immediate buffer against shocks. Land is unlikely to serve as collateral because the land market is very thin in the Kagera region of Tanzania, which also implies that reported values may be not be very reliable. Other types of physical assets are likely to be used as immediate buffers against shocks (e.g., cattle, which is a highly liquid asset). We examine the empirical validity of these claims below.

4. Results

We first examine whether crop shocks are plausibly exogenous and are economically significant. We then examine the effect of these shocks on child labor and whether the shocks have a smaller effect for households with a greater level of assets. Finally, we consider alternative interpretations of our results.

4.1. The occurrence and effect of crop shocks

We begin by examining empirically our claim in previous section that crop shocks are transitory and are a plausibly exogenous source of variation at the household level. The concern is that shocks, if persistent and recurrent in the same families over time, might pick up unobserved household characteristics rather than identifying an exogenous source of variation. However, as highlighted in Table 2, Panel A, the frequency of shocks in households over time is such that only 7% of the 483 families that were interviewed in all four rounds experienced shocks in three or four rounds. Table 2, Panel B, shows that shocks are substantial in magnitude: just under half the shocks wipe out between one and three

Table 2
Frequency and magnitude of shocks

Panel A: Frequency of shocks

Number of shocks across four survey rounds	Number of households	%
0	58	12.0
1	224	46.4
2	166	34.4
3	33	6.8
4	2	0.4
Total	483	100

Panel B: Magnitude of shocks, among households with one shock

Share of the value of crop loss to total crop value	% of sample
1–24%	38.8
25–49%	24.1
50–74%	25.0
75–100%	12.1
Observations	224

Table 3
Predicting the occurrence of shocks

	(1)	(2)	(3)
Dependent variable	Any crop loss	Any crop loss	Any crop loss
Specification	Probit	Probit	Probit
Lagged any crop loss		−0.004 (0.024)	−0.006 (0.024)
Lagged child labor intensity			0.001 (0.001)
Lagged log per capita assets	0.011 (0.007)	0.012 (0.007)	0.011 (0.007)
Head of the HH years of schooling	0.004 (0.004)	0.004 (0.004)	0.004 (0.004)
Head of the HH is female	0.049* (0.026)	0.048* (0.026)	0.048* (0.026)
Head of the HH's age/100	0.021 (0.070)	0.020 (0.070)	0.023 (0.070)
Observations	1585	1585	1585
R-squared	0.30	0.30	0.30

Regressions are at the household level. Marginal coefficients are presented. Standard errors are in parentheses. *** indicates significance at 1%; ** at 5%; and, * at 10%. Community, district, survey round, and season dummies are included but not reported.

quarters of the household's crop. Complete crop destruction is uncommon, which again is useful because complete crop destruction might render child labor useless.

In Table 3, we predict the occurrence of a shock based on current and lagged household characteristics, on a range of measures of household wealth, and on the lagged occurrence of shocks. In Table 3, column (1), we see that neither household characteristics nor lagged asset holdings are (individually or jointly) significant predictors of shocks, with the exception of female-headed household (which is a positive predictor of shocks). Column (2) introduces lagged crop loss. The coefficient is negative and not statistically significant, suggesting that agricultural shocks are transitory at the household level. The inclusion of household fixed effects in all of the subsequent specifications will further increase the credibility of this view. In column (3), we examine if households that use more child labor are more likely to experience shocks. Though this is plausible (and, if true in our data, would undermine a causal interpretation of the shock coefficient), we find no evidence that lagged child labor is a significant predictor of future shocks. Though this does not rule out all forms of reverse causation from child labor to shocks, it rules out the most obvious link and supports a causal interpretation of the crop shock effect.

Overall, we believe that these results lend credence to the use of crop shocks as a source of variation in our subsequent specifications.

4.2. The effect of crop shocks on child labor

Table 4 reports estimates of regressing child labor hours on crop shocks; all specifications control for wave and season fixed effects, in addition to age of the child. We begin with a household fixed effects specification that considers all children between the ages of 7 and 15 and controls for household characteristics. We find a positive and significant effect of effect of crop shocks. Children in households experiencing shocks increase their child labor by 6.1 h, or approximately 30%.

Table 4
Hours worked, income shocks, and assets

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable	Hours worked	Hours worked	Hours worked	Hours worked	Hours worked	Hours worked
Specification	FE	FE	FE with cluster-time trends	FE with cluster-time dummies	FE with cluster-time dummies	FE with cluster-time dummies
Sample restriction: land acres		1–25.5	1–25.5	1–25.5	1–25.5	1–25.5
Shock: any crop loss	6.08** (2.87)	9.04*** (3.24)	8.93*** (3.37)	7.92** (3.45)	7.61** (3.56)	7.50** (3.56)
Log per capita land value	0.47* (0.27)	0.77** (0.33)	0.80** (0.34)	0.91*** (0.35)	0.87** (0.35)	0.87** (0.35)
Shock x log per capita assets	−0.47* (0.28)	−0.78** (0.31)	−0.75** (0.32)	−0.66** (0.33)	−0.62* (0.35)	−0.61* (0.35)
Per capita land value					0.02*** (0.01)	0.02*** (0.01)
Shock/10 x per capita land value					−0.06 (0.23)	−0.06 (0.23)
Mother dead						−0.36 (0.86)
Father dead						0.63 (0.89)
Observations	5591	4709	4709	4709	4709	4706
R-squared	0.38	0.36	0.39	0.42	0.42	0.42

Standard errors are in parentheses. Standard errors are computed correcting for heteroskedasticity and correlation within household clusters. Fixed effects are at the household level. *** indicates significance at 1%; ** at 5%; and, * at 10%. Other regressors included, but omitted from the table, are mother's and father's education, age and age squared and indicator variables for missing parental education, the season at time of interview, and the round of the interview.

Although this is a reduced form result, and as such does not shed light on what mechanisms could give rise to the increase in child labor, it is of independent interest. It establishes that rural households are significantly exposed to crop shocks and that they respond to these shocks along a dimension that is undesirable from a policy perspective. It also suggests that insuring rural households against such shocks would be one means of reducing (although not eliminating) child labor.

There is a positive and significant relationship between the level of household assets and the use of child labor. A one standard deviation increase in log per capita assets leads to a 6% increase in child labor hours. As noted in the Introduction, though this is initially surprising (since child labor is often portrayed as being negatively associated with household wealth), in agricultural settings a positive association has been previously noted (see, inter alia, Bhalotra and Heady, 2003). Note that our measure of assets excludes land value, which would clearly have a positive association with the demand for child labor.

In the third row, we see that the effect of shocks is offset by household asset holdings. The interaction effect of assets and shocks is negative and significant. The mean level of assets offsets eighty percent of the effect of the crop shock. The ability of households to substantially offset the effect of shocks using household assets is consistent with a range of interpretations, including buffer stocks and borrowing. In the next section, we examine the evidence for these effects.

Before turning to issues of interpretation, we first subject our results to a series of robustness checks. In column (2), we exclude outliers from our sample, in particular landless and extremely wealthy households for whom crop shocks are unlikely to have a significant effect. For the restricted sample, the effect of shocks is substantially larger: a crop shock is associated with a 50% increase in child labor hours and the effect is significant at the one percent level. The positive effect of assets is also larger. The effect of the asset–shock interaction is now significant at the 5% level, and the mean level of assets offsets 87% of the effect of the shock.

In columns (3) and (4) we introduce cluster-time trends and cluster-time fixed effects, in addition to household fixed effects. Household fixed effects, in columns (1) and (2), control for time invariant household characteristics that could confound our results; among others these include a household's social network (which both acts as a buffer against shocks but could also be correlated with assets) and household agricultural production technology (which again is omitted and plausibly correlated with assets). Nonetheless, there is a range of time varying characteristics that remain a concern, including community level shocks to production technology or changes in the availability of schools. Community-time dummies control for these and other time varying shocks. Our coefficients of interest remain very robust in sign, significance, and magnitude.

In column (5) we examine whether controlling for land wealth affects our coefficients of interest. To the extent that land and non-land wealth are correlated, the coefficient on assets in columns (1) to (4) could be picking up the effect of land value; land values could mitigate the effect of shocks for the reasons we have discussed (buffer stocks and access to credit) but also for unrelated reasons (such as food self-sufficiency when faced with a shock). We see that land value is positively and significantly associated with child labor hours, but that it does not significantly offset the effect of a crop shock. Finally, in column (6), we control for parents' health status, which could both influence the crop shock and

independently lead children into child labor. Our coefficients of interest remain robust in sign, significance, and magnitude.

There are several alternative explanations for our results. The most direct concern is reverse causation: that the use of child labor causes shocks, rather than the reverse. This issue was examined in Section 4.1 (Table 3), where we showed that the use of child labor does not, in fact, predict the occurrence of shocks. More directly, evidence from the field interviews suggests that child labor on farms is used primarily in domestic activities relating to farming rather than direct agriculture such as planting or pest control. Thus, our evidence suggests that reverse causation is not a strong concern, though of course it cannot be ruled out.

Another concern is that shocks affect not only households' income but also their production technology. As such, a negative crop shock could directly induce an increased demand for child labor. We offer several pieces of evidence against this being the only explanation for our results. First, we have shown that the shock effect is weaker among households with a greater level of assets. Thus, the demand shock story would require that households with more assets also use a production technology that is less affected by crop shocks. Furthermore, to the extent that households are unlikely to switch production technology over the horizon covered by this survey (four waves, six months apart), we believe that our inclusion of household fixed effects renders this alternative explanation much less plausible. More generally, household, time, and cluster-time fixed effects rule out an array of omitted variable issues, such as village-level fertilizer use, pesticide price shocks, or household level tastes for child labor.

4.3. Extensions

In Table 5, we use our main specification (season and wave dummies, household fixed effects, and time-cluster dummies) to extend our results in several directions. In column

Table 5
Time allocation, income shocks, and collateral

	(1)	(2)	(3)	(4)
Dependent variable	Education	Communal farm hours	Firewood hours	Chore hours
Specification	FE with cluster-time dummies	FE with cluster-time dummies	FE with cluster-time dummies	FE with cluster-time dummies
Shock: any crop loss	−0.20** (0.09)	0.07 (0.11)	2.32*** (0.84)	4.94** (2.19)
Log per capita assets	−0.02*** (0.01)	0.01 (0.01)	0.16* (0.08)	0.69*** (0.22)
Shock x log per capita assets	0.02** (0.01)	−0.01 (0.01)	−0.20** (0.08)	−0.46** (0.21)
Observations	4709	4709	4709	4709
R-squared	0.57	0.34	0.35	0.42

Standard errors are in parentheses. Standard errors are computed correcting for heteroskedasticity and correlation within household clusters. Fixed effects are at the household level. *** indicates significance at 1%; ** at 5%; and, * at 10%. Other regressors included, but omitted from the table, are mother's and father's education, age and age squared and indicator variables for missing parental education, the season at time of interview, and the round of the interview.

(1) we examine whether crop shocks and asset holdings have a significant effect on education. Though it has been suggested that there is not a one-for-one trade-off between child labor and education (see Ravallion and Wodon, 2000), an effect on schooling would provide an important underpinning of our interest in child labor. We find that there is a negative and significant relationship between crop shocks and education. We find that children in households affected by a shock are twenty percentage points less likely to be enrolled (relatively to a mean enrolment rate of 70%). Log per capita assets are negatively associated with enrolment, though the magnitude of the effect is smaller (a one standard deviation increase in asset holdings leads to a 3% point decline in enrolment). Finally, looking at the interaction term, we find that the mean level of asset holdings almost fully offsets the negative enrolment effect of a shock.

We next investigate which components of child labor hours increase in response to shocks. In column (2) we note that time spent on communal farms is increasing in response to shocks and that this effect is smaller for wealthier households, but these effects are not statistically significant. In column (3) and (4) we find that child labor hours spent in gathering firewood and in other household chores both significantly increase in response to shocks. The number of hours spent gathering firewood almost doubles, and household chore hours increase by an order of 50%. For both activities, we find that households with a mean level of asset holdings are able to fully offset the effect of the shock.

The results in Table 5 demonstrate that not only do child labor hours increase in response to a shock, but that educational enrolment suffers. The increase in hours spent in activities such as chore hours and gathering firewood point to the reason for this decrease in enrolment.

4.4. Buffer stocks versus access to credit

In this section we investigate two alternative interpretations of our results. The fact that asset holdings substantially offset the effect of an agricultural shock on child labor hours is consistent both with a buffer stock story (households draw down their assets in times of need; see Deaton, 1992) and access to credit (households with a higher level of assets have greater access to credit and can borrow in times of need).

In Table 6, we first examine the effect of shocks on three categories of household assets: cash, physical assets (including land and business assets), and household durables. In the first row we find that there is a decrease in the holdings of all three assets in response to a shock (though only the effect on durables is significant). Cash holdings decline by more than 50% relative to the mean, holdings of physical assets decline by 15%, and durables decline by more than twice the value of mean holdings. Because the specification includes an interaction with overall assets, the results indicate that relatively poor households draw down their assets in response to a shock, consistent with a buffer stock view. However, the interaction effect with assets fully offsets the negative effect of shocks, suggesting that wealthier households are able to buffer their use of child labor and children's enrolment without drawing down asset holdings. This is consistent with wealthier households borrowing against their asset holdings.

Table 7 provides some direct evidence in this direction. We regress the probability of a household having an outstanding loan against shocks, asset holdings, and asset holdings

Table 6
The effect of crop shocks on household wealth

	(1)	(2)	(3)
Dependent variable	Per capita cash	Per capita physical assets	Per capita durables
Specification	FE with cluster-time dummies	FE with cluster-time dummies	FE with cluster-time dummies
Shock	−0.11 (0.79)	−24.05 (19.43)	−2.03** (0.88)
log per capita assets	−0.01 (0.09)	1.65 (2.23)	−0.09 (0.10)
Shock *x* log per capita assets	0.03 (0.08)	3.05 (1.88)	0.21** (0.09)
Observations	1462	1462	1462
R-squared	0.02	0.02	0.07

The specification in estimated at the household level. Regressions also control for education, sex, and age of household head, wave, season, and household fixed effects. Standard errors are in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

interacted with the shock. The direct effect of shocks is insignificant. This shows that low-wealth households are no more likely to take a loan in response to a shock, and is consistent with the buffer stock view of asset holdings suggested for low-wealth households in the context of Table 6. Not surprisingly, cash holdings are associated with a reduced probability of taking a loan, and the same is true in response to a shock. In contrast, households with durables are more likely to take a loan if they experience a shock. This corroborates the claim in Section 3 that durable assets are more likely to be used as collateral than physical assets (especially land), and is consistent with our discussion of Table 6, namely that higher wealth households appear to be more likely to take a loan in response to a shock rather than to draw down their assets.

These results are broadly consistent with the view that poorer households respond to shocks by drawing down their assets (in addition to increasing their use of child labor and disrupting their children's schooling) and that richer households instead resort to borrowing rather than drawing down their asset holdings. This does not immediately imply that households are credit constrained per se. In a world without moral hazard, it would be efficient for households — regardless of their asset holdings — to borrow when

Table 7
The probability of taking a loan in response to a shock

	Anyone in HH had a loan
Shock: any crop loss	0.022 (0.026)
Per capita cash holdings	−0.025* (0.015)
Shock *x* cash	−0.112* (0.065)
Per capita physical assets	−0.003 (0.178)
Shock *x* physical assets	−0.961 (0.672)
Per capita durables	−0.065 (0.172)
Shock *x* durables	1.604** (0.804)
Sample size	2042
Pseudo *R*-squared	0.096

Marginal probit coefficients are reported. The specification is estimated at the household level, and controls for household characteristics along with region and wave dummies.

faced with transitory shocks, rather than increase their use of child labor and disrupt their children's schooling, but this is not necessarily true in a world with moral hazard.

5. Conclusions

This paper has examined the link between crop shocks, household asset holdings, and child labor. Based on institutional knowledge of the Kagera region and on econometric analysis, we have argued that it is plausible to view crop shocks as transitory shocks to household income, that shocks are significant relative to household wealth, and that households use assets both as buffer stocks and as collateral for borrowing. We find that crop shocks lead to a significant increase in the level of child labor and that households with assets are able to offset approximately 80% of this shock. Educational enrolment decreases in response to shocks, but households with a typical level of asset holdings are able to fully offset the shock. Our results suggest that poorer households might be using assets as a buffer stock, drawing them down in times of need, whereas wealthier households' behavior is consistent with an access to credit story.

We have presented a range of evidence to corroborate this interpretation of our results: (i) our results control for a wide range of household characteristics and household fixed effects; this rules out time invariant (over a two year horizon) features of the household explaining away our results (e.g., the household's attitude toward child labor, its social network, or its production technology); (ii) our results are robust to controlling for cluster-time fixed effects; this rules out time-specific cluster shocks explaining away our results; (iii) we control for land wealth, and show that it is positively (though not robustly) associated with child labor and does not seem to have the buffering effect of other assets; and (iv) we show that households are more likely to take loans when they experience a shock and when they hold durable assets, and that the probability of taking a loan in response to a shock increases more among households holding durable assets.

It must, however, be acknowledged that there are alternative interpretations of our results. The effect of shocks on households could be explained by myopia or an extremely high discount rate relative to the interest rate. The differential effect of shocks on households with a high level of assets could be explained away if households with a higher level of assets also have a greater demand for credit when they experience a crop shock (although this would have to be true conditional on household fixed effects and other wealth, including landholding, for which we control). Although we cannot rule out all of these alternatives, we believe that we offer the most plausible interpretation of our results.

Our results have a number of important implications. Regardless of how they are interpreted, our results do demonstrate that households increase their use of child labor in response to crop shocks. To the extent that child labor is viewed as a policy problem, this result implies that policies that insure agricultural households against such shocks will lead to a reduced level of child labor. Furthermore, the credit interpretation of these results implies that increasing household access to credit, specifically in response to crop shocks, will reduce child labor and increase household welfare.

There are clearly many questions that we have not addressed in this paper. It would be ideal to examine loan-level information from agricultural households; this would fill in a

missing link between shocks and credit constraints. It is also important to consider whether the level of child labor that we observe in agricultural households in developing countries in fact has negative consequences on children's long-run development, as this is the impetus for efforts aimed at reducing child labor (see Beegle et al., 2004). These are subjects of ongoing research.

Acknowledgments

We thank Esther Duflo and the referees for their detailed comments, and Lena Edlund, Eric Edmonds, Ted Miguel, Martin Ravallion, John Strauss, Steven Zeldes, and seminar participants at the NEUDC 2002 conference, the World Bank Social Protection seminar series, the World Bank Economists' Forum, and the Center for the Study of African Economies 2004 conference at Oxford University for many useful suggestions. Support from the World Bank's Research Committee is gratefully acknowledged.

References

Ainsworth, M., Koda, G., Lwihula, G., Mujinja, P., Over, M., Semali, I., 1992. Measuring the Impact of Fatal Adult Illness in Sub-Saharan Africa. Living Standards Measurement Study, vol. 90. World Bank, Washington, D.C.

Ainsworth, M., Beegle, K., Koda, G., 2002. The Impact of Adult Mortality of Primary School Enrollment in Northwestern Tanzania. Africa Region Human Development Working Paper Series. World Bank.

Al-Samarrai, S., Peasgood, T., 1998. Education attainments and household characteristics in Tanzania. Economics of Education Review XVII, 395–417.

Al-Samarrai, S., Reilly, B., 2000. Urban and rural differences in primary school attendance: an empirical study for Tanzania. Journal of African Economies IX, 430–474.

Baland, J., Robinson, J., 2000. Is child labor inefficient? Journal of Political Economy CVII, 663–679.

Beegle, K., 2005. Labor effects of adult mortality in Tanzanian households. Economic Development and Cultural Change 53, 655–683.

Beegle, K., Dehejia, R., Gatti, R., 2004. Why Should We Care about Child Labor? The Education, Labor Market, and Health Consequences of Child Labor. National Bureau of Economic Research Working Paper, vol. 10980.

Bhalotra, S., Heady, C., 2003. Child farm labor: the wealth paradox. World Bank Economic Review 17, 197–227.

Bommier, A., Lambert, S., 2000. Education demand and age at school enrollment in Tanzania. Journal of Human Resources XXXV, 177–203.

Boozer, M., Suri, T., 2001. Child labor and schooling decisions in Ghana, mimeo.

Burke, K., 1998. Investing in children's human capital in Tanzania: does household illness play a role? Dissertation, SUNY Stony Brook, Stony Brook , NY.

Burke, K., Beegle, K., 2004. Why children aren't attending school: the case of rural Tanzania. Journal of African Economies 13, 333–355.

Canagarajah, S., Nielsen, H., 1999. Child labor and schooling in Africa: a comparative study, mimeo.

Chaudhuri, S., Ravallion, M., 1997. Risk and insurance in village India: comment. Econometrica 65, 171–184.

Deaton, A., 1992. Understanding Consumption. Clarendon Press, Oxford.

Dehejia, R., Gatti, R., in press. Child labor: the role of income variability and access to credit across countries, World Bank Policy Research Working Paper no. 2767 and National Bureau of Economic Research Working Paper No. 9018, Economic Development And Cultural Change.

Edmonds, E., 2002. Is child labor inefficient? Evidence from large cash transfers, mimeo.
Fallon, P., Tzannatos, Z., 1998. Child Labor: Issues and Directions for the World Bank. World Bank, Washington D.C.
Fitzsimmons, E., 2003. The Effect of Risk on Education and Child Labor. Institute for Fiscal Studies, Working Paper, vol. 02/07.
Foster, A., 1995. Prices, credit markets and child growth in low-income rural areas. Economic Journal 105 (430), 551–570.
Guarcello, L., Mealli, F., Rosati, F.C., 2003. Household Vulnerability and Child Labor: The Effect of Shocks, Credit Rationing and Insurance. Manuscript. UNICEF Innocenti Research Center.
Grootaert, C., Patrinos, H., 1999. The Policy Analysis of Child Labor: A Comparative Study. St. Martin's Press, New York.
Jacoby, H., 1994. Borrowing constraints and progress through school: evidence from Peru. The Review of Economics and Statistics LXXIV, 151–160.
Jacoby, H., Skoufias, E., 1997. Risk, financial markets, and human capital in a developing country. Review of Economic Studies LXIV, 311–335.
Kochar, A., 1999. Smoothing consumption by smoothing income: hours-of-work responses to idiosyncratic agricultural shocks in rural India. Review of Economics and Statistics LXXXI, 50–61.
Krueger, A., 1996. Observations on International Labor Standards and Trade. Working Paper, vol. 5632. National Bureau of Economic Research.
Mason, A., Khandker, S., 1997. Household schooling decisions in Tanzania. Draft, Poverty and Social Policy DepartmentWorld Bank, Washington, D.C.
Morduch, J., 1994. Poverty and vulnerability. American Economic Review Papers and Proceedings 84, 221–225.
Morduch, J., 1995. Income smoothing and consumption smoothing. Journal of Economic Perspectives IX, 103–114.
Morduch, J., 1999. Between the state and the market: can informal insurance patch the safety net. The World Bank Research Observer XIV, 187–207.
Ranjan, P., 2001. Credit constraints and the phenomenon of child labor. Journal of Development Economics LXIV, 81–102.
Ravallion, M., Wodon, Q., 2000. Does child labour displace schooling? Evidence on behavioral responses to an enrollment subsidy. The Economic Journal CX, 158–175.
Townsend, R., 1994. Risk and insurance in village India. Econometrica LXII, 539–591.
World Bank, 1993. Report of a workshop on: the economic impact of fatal illness in sub-Saharan Africa. World Bank, Washington, D.C. and the University of Dar es Salaam.
Zeldes, S., 1989. Consumption and liquidity constraints: an empirical investigation. Journal of Political Economy XCVII, 305–346.