Household Choices of Child Labor and Schooling

A Simple Model with Application to Brazil

Rodrigo R. Soares
Diana Kruger
Matias Berthelon

A B S T R A C T

This paper argues that conflicting results from previous literature—related to the effect of economic conditions on child labor—derive from different income and substitution effects implicit in different types of income variation. We use agricultural shocks to local economic activity in Brazil (coffee production) to distinguish between increases in household income and increases in the opportunity cost of time. Results show that higher household wealth is associated with lower child labor and higher schooling. Nevertheless, temporary increases in local economic activity are associated with higher child labor and lower schooling, particularly for children with poor economic backgrounds.

I. Introduction

This paper develops a simple theoretical model of household decisions regarding child labor and schooling. The model is used both as a benchmark

to read the available empirical evidence and as a guide to help in the choice of our empirical specification. Though both the theoretical and empirical literatures on child labor have blossomed in recent years, there remains a gap between the two. We argue that part of the conflicting results obtained in the empirical literature—related to the effect of improvements in economic conditions on child labor—derives from its lack of theoretical structure. Our theoretical model suggests an empirical specification where income and substitution effects from different components of family income can be clearly identified. We show that the effects of these components are different according to the margin of choice faced by the family, and incorporate these restrictions in the estimation. Our empirical results using Brazilian data show that family wealth tends to reduce the incidence of child labor and increase school attendance, while, conditional on family wealth, increases in economic activity (or in the opportunity cost of children's time) are associated with increased child labor and reduced school attendance. Most of the conflicting results from the previous empirical literature can be easily understood as a consequence of its incapacity to distinguish between the income and substitution effects implicit in different types of income variation.

Child labor has been identified as an important determinant of the persistence of poverty in developing countries. Inefficiently high levels of child labor lower human capital accumulation, reducing future wages and reproducing poverty and inequality. Baland and Robinson (2000) show that, even when socially inefficient, child labor may exist due to the incapacity of parents to borrow against the future income of children. Basu and Van (1998), in a multiple equilibria model, stress an alternative mechanism in which child labor is both a cause and a consequence of poverty: in a "good" equilibrium, when market wages are high, parents choose not to send their children to work; whereas in a "bad" equilibrium, when wages are low and families are poor, parents send their children into the labor force. Along similar lines, Dessy (2000) finds that there is a critical level of adult wages below which child labor is supplied.[1] Ranjan (2001) also shows that credit constraints lead to inefficiently high levels of child labor, which, in turn, are related to greater income inequality.

In these models, the market solution leads to a situation where it would be socially efficient to reduce children's labor supply and increase schooling. In such settings, child labor can be an intergenerational source of poverty traps: Poor children work today, obtain less schooling, and become low productivity workers, earning low wages in the future (as suggested by the evidence presented in Emerson and Souza 2003).[2]

Empirical work on the determinants of child labor has spanned a wide range of different settings and experiments. Edmonds (2005 and 2006), for example, finds

1. An exception in the theoretical literature is the model developed by Rogers and Swinnerton (2004), where low-income parents who anticipate future transfers from their children invest more in children's present education, whereas high-income parents do not expect or need such transfers and, therefore, invest less in the human capital of children. Though theoretically consistent, this model does not seem particularly relevant to analyze the current situation in most developing countries.
2. Although most of the empirical literature has focused on the relationship between child labor and school attainment or enrollment, Gunnarsson et al. (2006) find that child labor leads to lower standardized test scores among grade school children of nine Latin American countries, thus harming the quality of schooling as well.

that improvements in economic status explain most of the decline in child labor in Vietnam, while anticipated cash transfers to the elderly in South Africa are associated with increased schooling and reduced labor supply. Bourguignon et al. (2003) and Cardoso and Souza (2004) find that, in Brazil, conditional income transfers from the Bolsa Escola program increased the likelihood of schooling, but had no significant impact on the incidence of child labor.[3] Beegle et al. (2006) find that, in Tanzania, negative agricultural shocks (reports of value of crop losses due to insects, rodents, and other calamities) increase the number of hours worked by children and reduce school enrollment; they also show that households with a sufficiently high level of assets are able to fully offset the shocks. These results suggest that poverty and liquidity constraints are important determinants of household decisions regarding children's allocation of time.

But other empirical evidence reaches conclusions that may seem puzzling at first sight. Barros et al. (1994) find that, in the eight largest metropolitan areas of Brazil, child labor is higher during periods of low poverty and high economic growth, rather than during periods of economic downturns and high poverty. Similarly, also looking at urban Brazil, Neri and Thomas (2001) find that children are more likely to repeat a grade and more likely to work during periods of economic growth, while Duryea and Arends-Kuenning (2003) find that the incidence of child labor is higher and educational outcomes are worse when average wages are higher. Kruger (2006 and 2007) finds that, in coffee-producing regions of both Brazil and Nicaragua, children are more likely to work and less likely to go to school during periods of improved economic conditions due to coffee booms. Results such as these have led some authors—such as Barros et al. (1994) and Rogers and Swinnerton (2004)—to argue that there is at best a weak link between income and child labor, and to question whether poverty and credit constraints are indeed the sources of the problem.

We claim that the seemingly conflicting results from the empirical literature are entirely consistent with theory, once one realizes that different types of shocks to family income—or different types of experiments—bring together different combinations of income and substitution effects. Income changes that are mostly associated with changes in households' full income should represent either pure income effects or situations where income effects tend to be relatively more important, and therefore should increase the demand for schooling and reduce child labor. On the other hand, short-term fluctuations in wages, income, or economic growth—particularly when analyzed as deviations from a given secular trend—should be mostly associated with increases in the opportunity cost of children's time, given an expected present value of full income. Therefore, these changes should come close to isolating the substitution effect and should bring together increased child labor and reduced schooling.

In order to clarify how these different margins affect household choices regarding child labor and schooling, we develop a very simple model that captures the main empirical implications of the theoretical literature (Baland and Robinson 2000; Basu and Van 1998; Dessy 2000). In addition to making explicit the argument developed in the previous paragraph, the model suggests a specification that helps guide our empirical exercise. The vast majority of empirical work analyzes child labor and

---

3. Bolsa Escola is a public cash-transfer program conditional on children's school attendance.

schooling decisions as separate, discrete choices.[4] In our model, child labor and schooling are not mutually exclusive and the response of households to changes in income may vary with wealth. The theory leads to a formulation where, in order to account for all the relevant dimensions of the environment faced by families, child labor and schooling decisions are characterized as a generalized ordered discrete choice problem. This formulation highlights both the set of alternatives available to families (work only, schooling and work, and schooling only) and the fact that families choosing each one of these different alternatives may respond differently to marginal changes in exogenous variables. The theory also highlights the need to distinguish between changes in family full income and children's wages for the different dimensions of income and substitution effects to be adequately identified.[5]

Given the limitations of the data typically available, we are not able to estimate the structural model. Instead, we use it as a benchmark to guide our empirical specification and exploit shocks to coffee production in Brazil as exogenous variation in the demand for unskilled labor. We concentrate the analysis on Brazil's coffee-producing regions between 1993 and 2003, and use household data from the Brazilian Household Survey (Pesquisa Nacional por Amostras de Domicílio—PNAD) and coffee production data from the Brazilian Census Bureau (Instituto Brasileiro de Geografia e Estatística—IBGE). Controlling for family income and wealth, and for secular trends in income, we are able to distinguish between the effects of family income and increased demand for child labor (due to shocks to local economic activity). We find that, conditional on family socioeconomic status and wealth, and on long-term growth, exogenous shocks to local economic activity are associated with increased child labor and reduced schooling. Nevertheless, family socioeconomic status and wealth—as measured by hourly wage and job tenure of the head of the household, nonlabor income, and ownership of various durable goods—are associated with reduced child labor and increased schooling. These basic results hold even when we explore the more exogenous dimensions of variation in the value of coffee production by using municipality fixed-effects.

Additionally, in conducting various robustness checks, we assess the role of a series of demographic correlates of child labor. Older children are more likely to work, while children with higher educational attainment for a given age are more likely to be in school. Girls are more likely to be in school and not working, while children living in rural areas are more likely to be working and not in school. Conditional on household wealth and socioeconomic status, children living in households with many children are more likely to work and less likely to be in school.

The remainder of the paper is organized as follows. Section II presents a theoretical model of child labor and schooling. Section III describes the data used in the analysis and discusses the empirical specification. Section IV presents and discusses the results. Finally, Section V concludes the paper.

---

4. Exceptions are Levison et al. (2000) and Bourguignon et al. (2003), who estimate multinomial logits.
5. In the microsimulations performed by Bourguignon et al. (2003), the conditional cash transfer program analyzed has potentially both income and substitution effects (by reducing the opportunity cost of schooling and, simultaneously, increasing family income). But the authors do not address this issue explicitly, nor try to separately identify these two components.

II. A Simple Model of Child Labor and Schooling

A. The Household Problem

This section develops a simple model of the joint household decision of child labor and schooling. We concentrate on a simplified version of the household problem in order to obtain a solution that is as close as possible to an empirical specification. Though formally different, our model shares the same basic properties of the theory proposed by Basu and Van (1998).

Consider an economy where parents make all decisions. To keep things simple, we assume that households have only one parent and one child. The household's utility function is

$$(1) \quad U(c,h) = \frac{c^{\sigma}}{\sigma} + \beta h,$$

where c is household consumption, h is the human capital of the child, and β and σ are constant parameters, with $0 < \sigma < 1$ and $\beta > 0$.

Consumption goods are purchased with income from adult and child labor. Assume that parents participate fully in the labor market, so that their labor supply is fixed at the total amount of labor time available (t_p). Consumption satisfies the following budget constraint:

$$(2) \quad c \leq w_c l_c + w_p t_p,$$

where w_c is the child wage, w_p is the adult wage,[6] and l_c is the labor supply of the child. The child's time constraint is given by

$$(3) \quad l_c + e_c = t_c,$$

where e_c is the time spent on investments in human capital, and t_c is the total amount of time available for the child.

We abstract from the material costs of investments in education and assume that human capital is produced only with child's time, according to the technology $h = \alpha e_c \exp(v)$, where α is a technological constant and v is an individual specific factor. Several family and individual characteristics—ranging from education of parents to idiosyncratic ability or luck—may affect the productivity of investments in human capital. These are summarized in the term v. We model v as being a linear function of a vector of demographic characteristics of the household (x) and a random term (ε), such that $v = \gamma'x + \varepsilon$.

Writing down the full-income constraint and substituting for h, the household problem is

---

6. In principle, the adult wage depends on the level of human capital previously accumulated by the parent. We do not deal with this explicitly here, but it is the motivation behind the fact that the household attaches positive value to the human capital of the child (human capital is translated into higher future earnings).

(4) $\max\limits_{\{c,e_c\}}$ $\left\{\dfrac{c^\sigma}{\sigma}+\beta\alpha e_c\exp(v)\right\}$

subject to $c+w_c e_c \leq w_c t_c + w_p t_p.$

This very simple framework displays the same empirical implications stressed in the original work of Basu and Van (1998). First, child's and adult's labor supply are substitutes from the perspective of generating income for the household (*Substitution Axiom*). Second, since preferences are quasi-linear in the human capital of the child, the household's demand for child education respects what Basu and Van (1998) call the *Luxury Axiom*. In other words, for sufficiently low family consumption, the marginal utility of consumption is so high that the family allocates all of the child's time to the labor market. As family consumption grows and the marginal utility of consumption declines, the family eventually starts investing some of the child's time in acquiring education and, from this point on, additional family resources are devoted entirely to investments in the child's human capital. This remains true until the child dedicates her time fully to schooling, when again additional resources start being devoted to household consumption. The only difference from Basu and Van (1998) is that the minimal level of consumption, which in their case is given by what they call the subsistence level, here depends also on the wage that the child faces on the market. In other words, when deciding what is minimally acceptable to the family, parents are actually comparing the future benefits from the child's education with the current market value of the child's labor. Under these circumstances, the static nature of the model can be reinterpreted as describing an environment where families face credit constraints in their decisions about consumption and investments in children. Because both individual and intergenerational credit constraints are important issues for poor families in developing countries, we keep this interpretation throughout the paper.

Defining λ as the multiplier on the full-income constraint, first order conditions for c and e_c are, respectively,

(5) $c^{\sigma-1}=\lambda$, and

(6) $\alpha\beta\exp(v)\lesseqqgtr\lambda w_c,$

with the inequality holding as $<$ when $e_c=0$, as $=$ when $0<e_c<t_c$, and as $>$ when $e_c=t_c$. Substituting for λ in the second expression, we characterize the household decision. If $\alpha\beta\exp(v)<c^{\sigma-1}w_c$, the marginal gain from investments in human capital is lower than the marginal value of child labor, so the child works and does not go to school. If $\alpha\beta\exp(v)=c^{\sigma-1}w_c$, the marginal value of one unit of time invested in human capital equals the marginal value of one unit of time supplied in the labor market, so the child shares her time between work and schooling. If $\alpha\beta\exp(v)>c^{\sigma-1}w_c$, the marginal value of one unit of time invested in human capital is higher, so the child spends all her time studying.[7] Each of the three choices available to families are discussed and characterized in further detail below.

---

7. A fully recursive model, where parents utility depended on the child's utility, which in turn would depend on wages, would lead to a similar characterization of the three alternative choices available to

1. First Case: Work and No Schooling

This choice is characterized by the inequality $\alpha\beta\exp(v) < c^{\sigma-1}w_c$. In this case, all the child's time is used as labor supply, such that $l_c = t_c$ and $e_c = 0$. Families in this situation already send their children to work, and, for small changes in the economy, will never switch to a situation where the children go to school and are entirely out of the labor market. What they contemplate is the possibility of starting to invest part of the children's time on schooling. From the budget constraint, consumption is given by $c = w_c t_c + w_p t_p$, so that $\alpha\beta\exp(v) < (w_c t_c + w_p t_p)^{\sigma-1}w_c$. Substituting the expression for v and using natural logarithms:

$$(7) \quad \varepsilon < \ln(1/\alpha\beta) + \ln w_c + (\sigma-1)\ln(w_c t_c + w_p t_p) - \gamma'x.$$

This formulation decomposes the income and substitution effects and illustrates the forces that will also play a role for families in other margins of choice. The $\ln w_c$ term represents the pure substitution effect, while family full income ($\ln(w_c t_c + w_p t_p)$) represents the pure income effect. For constant family income, a higher wage for the child is associated with a higher opportunity cost of schooling and, therefore, a higher probability that the child will work and not go to school. For constant wage of the child, higher family income is associated with a lower probability that the child will work and not go to school. An uncompensated change in w_c, however, does not have an unequivocal effect.

2. Second Case: Schooling and No Work

This choice is characterized by the inequality $\alpha\beta\exp(v) > c^{\sigma-1}w_c$. In this case, all the child's time is used on investments in human capital, so that $l_c = 0$ and $e_c = t_c$. From the budget constraint, consumption is given by $c = w_p t_p$, so that the inequality can be written as $\alpha\beta\exp(v) > w_c(w_p t_p)^{\sigma-1}$. Substituting the expression for v and using natural logarithms:

$$(8) \quad \varepsilon \geq (1/\alpha\beta) + \ln w_c + (\sigma-1)\ln(w_p t_p) - \gamma'x.$$

In this case, the child does not work, so the income effect is captured only by the full income of parents ($\ln(w_p t_p)$). Marginal changes in children's wages do not affect the consumption of the family and, therefore, have no income effect, but they do affect the opportunity cost of investments in human capital.

3. Third Case: Work and Schooling

This is the intermediary case, with $\alpha\beta\exp(v) = c^{\sigma-1}w_c$. In this situation, we have $l_c > 0$ and $e_c > 0$. In words, the children share their time between work and schooling. From the first order conditions, consumption is given by $c = [\alpha\beta\exp(v)/w_c]^{1/(\sigma-1)}$.

---

families. The only difference in this case would be that the term $c^{\sigma-1}$ would be replaced by the inverse of the growth rate of consumption across generations, and another constant term (t_p) would appear multiplying the left-hand side of the expression. Still, the basic trade-off would be between the future welfare of the child (in our model represented directly via a reduced form by h) and the current welfare of the family. We keep the formulation in the text because it is simpler and closer to what an empirical specification would look like.

Using the expression for v, this situation is characterized by the following inequalities

(9) $\ln(1/\alpha\beta) + \ln w_c + (\sigma - 1)\ln(w_c t_c + w_p t_p) - \gamma' x \leq \varepsilon \leq \ln(1/\alpha\beta)$
$$+ \ln w_c + (\sigma - 1)\ln(w_p t_p) - \gamma' x.$$

B. Extensive Margin Choice and Empirical Specification

The household faces a discrete choice with three possible options in relation to the child: work and no schooling, work and schooling, and schooling and no work. Define the discrete variable J indicating the household choice as 0, 1, and 2, accordingly. The household decision is

$$J = \begin{cases} 0, \text{ if } \varepsilon < 1/\alpha\beta + \ln w_c + (\sigma - 1)\ln(w_p t_p + w_c t_c) - \gamma' x, \\ 1, \text{ if } 1/\alpha\beta + \ln w_c + (\sigma - 1)\ln(w_p t_p + w_c t_c) - \gamma' x \leq \varepsilon \\ \qquad\qquad < 1/\alpha\beta + \ln w_c + (\sigma - 1)\ln(w_p t_p) - \gamma' x, \\ 2, \text{ if } \varepsilon \geq 1/\alpha\beta + \ln w_c + (\sigma - 1)\ln(w_p t_p) - \gamma' x. \end{cases}$$

This is a generalized ordered discrete choice model. The realization of the random variable ε determines the choice of the household. The difference from standard ordered models lies in the fact that the threshold point is not constant. Here it is given by the explanatory variables, implicit in the difference between $\ln(w_p t_p + w_c t_c)$ and $\ln(w_p t_p)$.

If all variables discussed above were observable, this simple model would lead immediately to an empirical specification. Consider a sample of individuals and index observations related to the ith child by the subscript i, with $i \in I = \{1, \ldots, n\}$, where n is the number of observations in the sample. Assume that ε is randomly distributed with distribution function $F(.)$. Additionally, define $Z_{ij} = 1$ if child i is in state j, with $Z_{ij} = 0$ otherwise, where $j \in J = \{0, 1, 2\}$. The likelihood function for the household problem is

(10) $$L = \prod_{i=1}^{n} \left\{ \begin{array}{l} [F(1/\alpha\beta + \ln w_{ci} + (\sigma - 1)\ln(w_{pi} t_p + w_{ci} t_c) - \gamma' x_i)]^{Z_{i0}} \times \\ \left[\begin{array}{l} F(1/\alpha\beta + \ln w_{ci} + (\sigma - 1)\ln(w_{pi} t_p) - \gamma' x_i) - \\ \quad - F(1/\alpha\beta + \ln w_{ci} + (\sigma - 1)\ln(w_{pi} t_p + w_{ci} t_c) - \gamma' x_i) \end{array} \right]^{Z_{i1}} \times \\ [1 - F(1/\alpha\beta + \ln w_{ci} + (\sigma - 1)\ln(w_{pi} t_p) - \gamma' x_i)]^{Z_{i2}}. \end{array} \right\}$$

But we cannot estimate this model structurally for two reasons, in addition to the overly simplifying assumptions implicit in the theoretical formulation. First, for the vast majority of children, we do not observe market wages, and market wages may be endogenous to child labor decisions in previous periods. And second, we do not observe full family income.

In relation to child wages, we choose to use a proxy for the demand for child labor at the municipality where the child lives. This allows us to use all the children in the sample, including the large number of those who do not work. The proxy

used is the value of coffee production per capita, which is discussed in detail later on. The basic idea is that this variable is correlated with the local demand for unskilled labor, and at the same time have some degree of exogenous variation due to the uncertainty associated with climatic conditions and agricultural production.

In relation to the second problem—nonobservable full income—we adopt the following strategy to motivate our empirical strategy. The logarithm of household full income can be rewritten as $\ln(w_{ci}t_c + w_{pi}t_p) = \ln w_{pi} + \ln t_p - \ln s_{pi}$, where $s_{pi} = t_p w_{pi}/(t_p w_{pi} + t_c w_{ci})$. The term s_{pi} gives the share of the family's full income that comes from parents, or, alternatively, the relative importance of the child in terms of the income generating power of the household. Notice that it refers to full income, so that, given the educational level of parents and the market wages, it does not depend on household decisions at that specific point in time. It will typically depend on family characteristics such as educational attainment and age of the different members, gender of the child, composition of the household, wealth, etc. Since $0 \leq s_{pi} \leq 1$, we have $\ln s_{pi} \leq 0$, and the discrete choice problem discussed before can be rewritten as

$$(11) \quad J = \begin{cases} 0, \text{ if } \varepsilon_i < \alpha^* + \ln w_{ci} + (\sigma - 1)\ln w_{pi} - \gamma' x_i - (\sigma - 1)\ln s_{pi}, \\ 1, \text{ if } \alpha^* + \ln w_{ci} + (\sigma - 1)\ln w_{pi} - \gamma' x_i - (\sigma - 1)\ln s_{pi} \\ \qquad\qquad\qquad \leq \varepsilon_i < \alpha^* + \ln w_{ci} + (\sigma - 1)\ln w_{pi} - \gamma' x_i, \\ 2, \text{ if } \varepsilon_i \geq \alpha^* + \ln w_c i + (\sigma - 1)\ln w_{pi} - \gamma' x_i, \end{cases}$$

where $\alpha^* = \ln(t_p^{\sigma-1}/\alpha\beta)$. Given some distribution function $F(.)$, the likelihood function for this problem is analogous to the one in Equation 10.

Our actual empirical implementation is akin to assuming that s_{pi} can be written as a function of a set of demographic characteristics z_i, as $\ln s_{pi} = \theta' z_i$.[8] We save the discussion on the components of the vectors x_i and z_i until the next section.

The formulation with these simplifying assumptions, which does not constitute a structural estimation of the model but draws heavily from it, constitutes our benchmark specification. Other empirical issues arising from the limitations of the theory (households with only one child, parents employed full time, etc) and from the proxy used for the demand for child labor are dealt with in the empirical section.

III. Data and Variables

A. Data Sources

The data used come from four different sources. All household variables are constructed from nine rounds of the Brazilian Household Survey (PNAD), which contains information on characteristics of all household members. The PNAD is con-

---

8. In principle, z_i and x_i can contain different elements or can be the same vector. If they contain only different elements, the variables in z_i will be responsible for determining the shift in the threshold as family characteristics change. Alternatively, if z_i and x_i are the same vectors, the threshold will be identified by allowing the coefficients on the variables in x_i to vary across the limits of the different categories.

ducted nationally throughout Brazil during the month of September of each year. We restrict the analysis to the period between 1993 and 2003, for which consistent sampling methodologies and questionnaires were maintained.[9]

The second source of data is the municipality-level surveys of agricultural production administered each year by the Brazilian Census Bureau (Instituto Brasileiro de Geografia e Estatística—IBGE), which we use to construct the value of coffee production per capita. Third, we use municipality level GDP and population data from the Brazilian Institute of Applied Economic Research (Instituto de Pesquisa Econômica Aplicada—IPEA) and from IBGE. Finally, at a later stage, we use yearly data on temperature and rainfall by quarter (December to February, March to May, June to August, and September to November). The quarterly data are originally from the Climate Research Unit of the University of East Anglia.

B. Variables and Sample Selection

Our dependent variable follows the empirical specification laid out by Equation 11 and reflects the child's schoolwork decision ordered according to what may be regarded as increasingly "better" outcomes.[10] It is equal to zero if the child works exclusively, one if the child works and goes to school, and two if she only goes to school (the definition refers only to market work, excluding household chores). The variable is constructed based on two questions asking whether the child worked last week or at any point during the past 12 months, and whether she is currently enrolled in school.[11] Our definition of child work is not restrictive and includes children who worked at any point during the previous 12 months.

Three sets of independent variables constitute our main interest. First, we need variables that capture wage variations for both parents (w_{pi}) and children (w_{ci}). Second, we need a set of variables capturing the return to schooling for each particular child (x_i). And, finally, we need a set of variables indicating the relative importance of the child's income in terms of the full income of the household (z_i).

1. Wages and Proxies for the Demand for Child Labor

For the first set of variables we use as parent's hourly wage the head of the household hourly wage, constructed from data on labor earnings and hours worked. In order to minimize the problem of interaction between child labor and labor supply decisions of adults, and to bring the sample closer to our theoretical model, we restrict the analysis to observations where the head of the household is employed full time.

---

9. We include the following years: 1993, 1995–1999, 2001–2003. In 1994 and 2000 PNAD was not conducted.

10. This statement is not entirely precise, since a family can always be made better off with the child working (in comparison to the child not working), as long as the child's wage is sufficiently high. Nevertheless, these choices are indeed ordered from the perspective of the family choice (the statement would be precisely correct if we thought in terms of wealth or compensated changes in full income).

11. A fourth outcome would be if the child responded no to both questions, which could be interpreted as dedicating time exclusively to leisure. However, we believe that children are not likely to be completely idle, so that this response probably hides some type of work—domestic or in the market—or a transitory state. Rather than reassigning this outcome to one of the other groups, we do not include it in the empirical analysis; less than 4 percent of children in the sample fall in this category.

Later on we also introduce additional variables to control for the interaction between labor supply of children and labor supply of other adults in the household. We define working full time for an adult as working at least 30 hours per week. Then we define the head of the household as the spouse with the highest hourly wage (in case both spouses work full time). We also restrict the sample to children aged 10 to 14, to heads of the household between 18 and 65 years of age, and to children who are relatives of the head of the household, in order to concentrate the analysis on a more homogeneous group.

As mentioned in the previous section, we do not observe child wages for most children, so we need to use proxies for the level of demand for child labor. We use the value of coffee production per capita (in logarithms) as a proxy for the level of economic activity, which may affect the opportunity cost of children's time. The variable is constructed for the municipality where the child lives.

We see the value of coffee production per capita as a proxy for the demand for low-skilled labor, for which child labor could potentially be a substitute.[12] In municipalities where coffee is economically important, the value of coffee production is an indicator of local economic activity and, at the same time, retains a certain level of exogeneity, given the uncertainties generally associated with climatic conditions and agricultural production. Similar identification strategies were used by Schultz (1985), where commodity prices served as instruments for the opportunity cost of women's time, by Black et al. (2003), where the price of coal was used as an instrument for men's wages, and by Kruger (2006 and 2007), where coffee production was used as a proxy for the value of children's time.

The underlying identifying assumption here is that, conditional on controls, the wage of the head of the household and per capita coffee production are uncorrelated with unobservables. This is needed for the estimated coefficients to reflect the causal effect of these variables on children's allocation of time. In order to deal with concerns that coffee production may be correlated with municipality characteristics, as when poorer municipalities have more agricultural production and more child labor, we also include municipality dummies in most of the results presented. This specification goes one step further in isolating shocks to local economic activity since it only requires that—conditional on controls, on the wage of the head of the household, and on constant municipality characteristics—variations in coffee production over time are not correlated with the error term.

In a final robustness exercise, we instrument the value of coffee production per capita with climatic variables (ignoring the problems related to instrumental variables in nonlinear contexts discussed, for example, by Angrist and Pischke 2009; the specific way in which we implement this procedure is outlined in detail in the results section). In this particular case, the identifying assumptions would be that, conditional on controls and on the wage of the head of the household, rain and temperature are uncorrelated with unobservables and do not directly affect child schooling and labor choices.

---

12. Parents' wages are also probably affected by coffee production. But since our specification explicitly includes parents' wage as an additional independent variable, our coffee-production variable will be capturing the effect on the demand for child labor orthogonal to that coming from its impact on parental wages.

Table 1
*Distribution of Children by Year (percent), Brazil, 1993–2003*

Year	Work Only	Work and School	School Only	Observations
1993	11.50	22.87	65.63	3,157
1995	6.95	21.66	71.39	3,223
1996	4.50	17.01	78.49	2,892
1997	3.79	17.85	78.36	3,036
1998	3.56	17.27	79.17	2,919
1999	2.74	18.34	78.92	2,993
2001	1.77	10.91	87.32	2,714
2002	1.30	12.73	85.97	2,773
2003	0.71	11.34	87.95	2,664
Total	4.25	16.9	78.84	26,371

Notes: The sample is restricted to children aged between 10 and 14 related to head of the household (heads of household are restricted to fully employed and aged between 18 and 65), and to municipalities with fewer than 100,000 inhabitants in the top 60 percent of coffee-producing municipalities. Data are from PNAD and IBGE agricultural surveys.

Finally, we concentrate the analysis on municipalities where coffee is an important economic activity. We see increased demand for child labor not as being necessarily linked to harvesting or actual agricultural production. Where coffee is an important activity, increased production may increase demand directly through demand for work in agriculture, but also indirectly, through transportation, processing, and packaging of coffee, as well as other auxiliary and satellite activities affected by coffee cycles. So in order to concentrate the analysis on localities where coffee production is indeed important, we restrict the sample to municipalities in the top 60 percent of coffee-producing municipalities. In addition, to concentrate on more homogeneous localities, we look only at municipalities with total population equal to less than 100,000 inhabitants.

The incidence of the three possible outcomes for each year is presented in Table 1. Figure 1 also presents the distribution of weekly hours of work for children included in the sample. As can be seen from the table, work only exhibits a sharply declining trend with the incidence falling from more than 10 percent to less than 1 percent between 1993 and 2003. Work and school shows a clear declining trend, from 23 percent in 1993 to 11 percent in 2003, while the school only outcome exhibits an increasing trend. Figure 1 shows that roughly 85 percent of working children in our sample worked at least 10 hours per week at some point during the previous 12 months, and at least 75 percent worked more than 15 hours per week.

Table 2 presents the value of coffee production per capita for municipalities included in the sample using the selection described before. The value of coffee production per capita increases steadily between 1993 and 1999. In 2001, it collapses—along with the international and domestic prices of coffee—to 55 percent of the
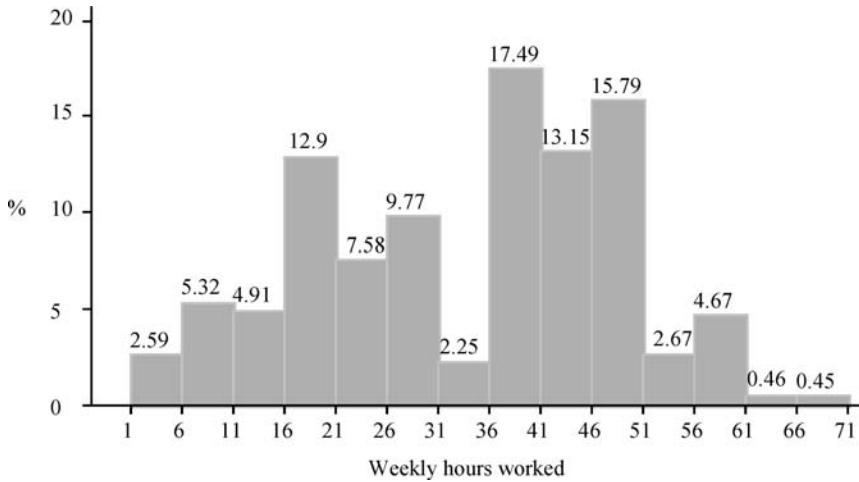
Figure 1
Distribution of Weekly Hours Worked, Children Included in the Sample, 1993–2003

Table 2
Yearly Coffee Production per capita, Brazilian Municipalities included in the Sample, 1993–2003

Year	Average Production Per Capita	Number of Municipalities
1993	107.29	142
1995	150.64	143
1996	157.05	142
1997	194.14	142
1998	206.56	141
1999	213.14	140
2001	94.55	140
2002	125.54	139
2003	123.09	139

Notes: Production per capita is in 2000 R$ (reais). The sample is restricted to children aged between 10 and 14 related to head of the household (heads of household are restricted to fully employed and aged between 18 and 65), and to municipalities with fewer than 100,000 inhabitants in the top 60 percent of coffee-producing municipalities. Data are from PNAD and IBGE agricultural surveys.

1999 value and below the level observed at the beginning of the period. Afterward, production experiences a recovery, but still remains below the 1995 levels. Our sample includes roughly 140 municipalities throughout the period.

2. Other Independent Variables and Additional Controls

In relation to the variables affecting the return to schooling (x_i), we include years of schooling, gender, race, and the age of the child, along with whether the household lives in a rural area.We also include the age of the head of the household, years of education of the head of the household, and whether the head of the household is female. The age, gender and race of the child, previous investments in human capital, and rural location may all affect the marginal gain from additional years of education. The age and education of the household head, and whether the household is headed by a woman, may reflect family characteristics that can be seen as direct inputs into the human capital production function and, because of that, may affect the productivity of investments in education.

The importance of the child's earning potential in relation to the household's full income (z_i) depends, to a great extent, on a set of variables very similar to that determining the return to schooling. In this respect, we want variables capturing the child's earning potential and also indicators of the household's full income (or wealth). The child's age, educational attainment, and the location of the household in a rural area may all affect her earning potential. Parent's tenure in the current job, other income of the household, and variables related to household infrastructure are all closely related to the household's full income and wealth. In relation to household infrastructure, we consider variables that indicate socioeconomic status, such as the number of bedrooms per person, availability of electricity and ownership of a telephone line (fixed), a television set, a refrigerator, and a washing machine. These are durable goods typically ordered in terms of family choices, so that they are closer to indicating differences in socioeconomic status, rather than differences in tastes.

Variables included only in x_i will appear with the same coefficient in the first and second transition equations shown in Equation 11. Variables included in both x_i and z_i will also appear in both equations, but with different coefficients in each one. Finally, variables included in z_i, but not in x_i, will appear only in the first equation, denoting the transition between state 0 (child labor and no schooling) and state 1 (child labor and schooling).

Because our theoretical model is overly simplified and seen only as a guide to our empirical strategy, after this initial specification we adopt a more flexible approach and relax some of the cross-equation constraints imposed by the theory. This less rigid specification is then used to conduct a series of robustness tests. One important concern is related to factors not included in the discussion above, but that may be correlated with the explanatory variables and with household decisions. As mentioned before, to account for geographic variables that may be correlated with coffee production and child labor, we use municipality fixed effects in most specifications. To account for the role of housework, intra-household substitution of labor, and dilution of family resources across different children, we control for: the presence of other children in the household, identified by particular age groups (age five and younger, between six and nine, and between 15 and 18); the total number of siblings; the presence of a person above 60 years of age in the household; whether both parents work; and whether the head of the household is a single parent. Also, in order to control for variations in local characteristics, our robustness checks in-

clude as additional independent variables municipality GDP per capita, population, and state-specific time dummies. We also test the robustness of the results to several different restrictions on the sample, to alternative definitions of the dependent variable and, in a final exercise, to the use of instruments for the value of coffee production per capita (based on climatic variables). Detailed discussion on these alternative specifications is saved until the next section. Table 3 presents summary statistics for the explanatory variables.

IV. Results

We assume that the error term ε follows a logistic distribution, so that the choice problem described in Equation 11 can be represented by a generalized ordered logit model.[13] Again, the dependent variable takes on the following three values: 0 if the child did not attend school and worked at any time during the previous 12 months, 1 if the child worked and went to school simultaneously, and 2 if the child attended school and did not work during the previous year. The generalized aspect of the model means that the proportional odds assumption is not maintained, so that thresholds between the three outcomes may vary according to families' characteristics.

A. Baseline Results

Table 4 presents the initial results of our estimation. The first two columns contain the estimated parameters for the two transition equations under our initial specification. This specification includes the value of coffee production per capita, the wage of the head of the household (both constrained to have the same coefficients across the two equations), and child and household characteristics, which are assumed to be part of x_i and z_i so that they are included in both transition equations, but with different coefficients in each one (not shown in the table). This initial specification also includes year and state dummies. Implicitly,we are assuming that variables affecting the child's return to schooling also affect the relative importance of her income to the household. In the third and forth columns, we present the results departing from the strict implications of the simplified theoretical model and allowing the coefficients on coffee production per capita and wage of the head of the household to vary across equations. In all following results, we adopt this more flexible specification and allow the coefficients to vary across equations. In the fifth and sixth columns, we include the socioeconomic variables as additional controls. As these variables are related to the importance of the child's income to the household, they are part of z_i but not of x_i, and therefore appear only in the first transition equation. Finally, in the last two columns we drop the state dummies and include municipality dummies.

---

13. We assume that the error term follows a logistic rather than a normal distribution because the estimation of generalized ordered probit models is particularly cumbersome from the numerical perspective. Our estimation was implemented using the commands and guidelines discussed in Williams (2006).

Table 3

Summary Statistics, Brazil, 1993–2003

Variable	Observations	Mean	Standard Deviation	Minimum	Maximum
Income variables					
Value of coffee production per capita	26,371	95.8	200.5	0	2,343
Hourly wage of head of household	26,371	1.6	2.8	0	172
Child characteristics					
Years of education	26,371	4.4	1.9	0	17
Female	26,371	0.49	0.5	0	1
Mixed	26,371	0.43	0.5	0	1
Black	26,371	0.05	0.21	0	1
Age	26,371	12.0	1.4	10	14
Lives in rural area	26,371	0.29	0.46	0	1
Family characteristics					
Age of head of household	26,371	41.7	8.0	18	65
Education of head of household (years)	26,371	5.5	4.2	0	17
Female head of household	26,371	0.25	0.44	0	1
Wealth characteristics					
Bedrooms per person	26,336	0.49	0.16	0.07	1.50
Has electricity	26,337	0.92	0.27	0	1
Has telephone	26,326	0.26	0.44	0	1
Has television	26,313	0.83	0.38	0	1
Has refrigerator	26,336	0.75	0.43	0	1
Has washing machine	26,339	0.21	0.41	0	1
Tenure of head of household (months)	26,371	109	109	0	696
Other income of household	26,371	36	149	1	7,025
Family structure					
Both parents work	26,371	0.32	0.47	0	1
Head of the household is single parent	26,371	0.09	0.29	0	1
Siblings 0–5 years living in household	26,371	0.30	0.46	0	1
Siblings 6–9 years living in household	26,371	0.42	0.49	0	1
Siblings 15–18 years living in household	26,371	0.39	0.49	0	1
Total number of siblings	26,371	1.43	1.28	0	9
Elderly member present	26,371	0.04	0.19	0	1

Notes: Monetary values are in 2000 R$ (reais). The sample is restricted to children aged between 10 and 14 related to the head of the household (heads of household are restricted to fully employed and aged between 18 and 65), and to municipalities with fewer than 100,000 inhabitants in the top 60 percent of coffee-producing municipalities. Data are from PNAD and IBGE agricultural surveys.

Table 4
Household Choices of Child Labor and Schooling—Generalized Ordered Logits, Brazil, 1993–2003

Coefficient	Model 1 Work/Work and School	Model 1 Work and School/School	Model 2 Work/Work and School	Model 2 Work and School/School	Model 3 Work/Work and School	Model 3 Work and School/School	Model 4 Work/Work and School	Model 4 Work and School/School
Coffee value (ln)	−0.0450***	−0.0450***	−0.0800***	−0.0405***	−0.0724***	−0.0406***	−0.0924***	−0.0163
	[0.0142]	[0.0142]	[0.0156]	[0.0153]	[0.0150]	[0.0153]	[0.0334]	[0.0242]
Household wealth (ln)	0.298***	0.298***	0.295***	0.299***	0.217***	0.302***	0.167**	0.249***
	[0.0337]	[0.0337]	[0.0621]	[0.0349]	[0.0621]	[0.0346]	[0.0739]	[0.0370]
Demographic characteristics	X	X	X	X	X	X	X	X
Household wealth and infrastructure							X	X
Year fixed effects	X	X	X	X	X	X	X	X
State fixed effects	X	X	X	X	X	X		
Municipal fixed effects							X	X
Constant	12.38***	6.468***	12.24***	6.487***	12.03***	6.485***	12.10***	6.160***
	[0.497]	[0.274]	[0.508]	[0.275]	[0.545]	[0.277]	[0.590]	[0.278]
Number of observations	26,371		26,371		26,293		26,293	
Number of municipalities	143		143		143		143	
Chi-squared	15,377				19,120		3.04E+11	
Chi-squared p-value	0.00				0.00		0.00	

Notes: State (or municipality) and year fixed effects are included in all specifications. Robust standard errors are clustered at the municipality level in parentheses. *** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$. Unit of observation is child. Dependent variable is categorical, indicating whether child works only (0), works and goes to school (1), or goes to school only (2). Independent variables are the value of the municipality's coffee or agricultural production per capita (ln), the hourly wage of head of household (ln), demographic characteristics (education, female dummy, race dummies, age, rural dummy, age of head of household, education of head of household, and female head of household dummy), household wealth and infrastructure (number of bedrooms per capita, dummies indicating whether household has electricity, phone, television, refrigerator, and washing machine, tenure of head of household in current job, and ln of other household income), year dummies, and location dummies (state or municipality). Results on variables not shown in the table are available upon request. A model from the first two columns imposes the constraint of equal coefficients on the wages of the household head and coffee production across equations. The sample is restricted to children aged between 10 and 14 related to head of the household (heads of household are restricted to fully employed and aged between 18 and 65), and to municipalities with fewer than 100,000 inhabitants in top 60 percent of coffee-producing municipalities. Data are from PNAD and IBGE agricultural surveys.

The coefficients from Table 4 refer to the effects of the explanatory variables on the likelihood of working only compared to going to school and working (Columns 1, 3, 5, and 7), and on the likelihood of going to school and working compared to going to school only (Columns 2, 4, 6, and 8). The estimated coefficients measure the effect of the independent variables on the likelihood of higher-valued outcomes. So a positive and significant coefficient on the first transition equation means that increases in the independent variable are associated, for families within that margin of choice, with a higher likelihood that the child works and goes to school, instead of only working (and analogously for the second transition equation).

In the first two columns, we see that conditional on the hourly wage of the head of the household, higher values of coffee production per capita are associated with worse outcomes for children. At the same time, conditional on the value of coffee production, higher hourly wages for the head of the household are associated with better outcomes for children. The two following models show that the same qualitative pattern persists when the coefficients on the variables of interest are allowed to vary across equations, or when the household infrastructure variables are included as controls for socioeconomic background. The only noticeable difference is that the effect of changes in coffee production becomes particularly strong for children on the margin of choice between work only and work and schooling.

This pattern is intensified once we include municipality dummies. In the last two columns, the coefficient on the value of coffee production for children in the margin of dropping out of school becomes even larger in magnitude, while that for children on the margin between schooling only and work and school becomes smaller in magnitude and not statistically significant. Overall, the results from Table 4 suggest that, conditioning on the hourly wage of the head of the household, on determinants of the return to schooling and on the relative importance of the child's income for the family, as well as on year- and location-specific dummies, the value of coffee production per capita highlights the substitution effect from increased economic activity.[14] The evidence suggests an intuitively appealing pattern where the effect of the shock to local economic activity represented by the value of coffee production is more important for poorer families, who are on the margin of deciding whether to send a child who works to school. Given the type of variation in labor demand captured by our coffee variable, this seems to be expected.[15]

Most of the control variables also have the expected effects (not shown in the table, but available from the authors upon request). Children who are older or live in rural areas are more likely to work and less likely to be in school, while, conditional on age, children with more years of schooling are more likely to stay in school. Also, children with more educated parents are more likely to be in school

---

14. In reality, increases in demand for child labor also always bring together a bit of income and substitution effects. But the results show that our strategy, and the controls that we include in the estimation, tend to isolate the substitution effect. This result is consistent with that in Kruger (2007).

15. For the interested reader, and to illustrate the advantages of using our generalized ordered discrete choice model, Appendix Table A1 presents results analogous to the ones from Table 4, but estimated with simple logits (individual level data) or linear models (data aggregated at the municipality level). None of the models reveal the patterns that appear in Table 4. These models are unable to deal with the differential responses that are at the core of our proposed strategy.

and not to work, while children in families with better socioeconomic characteristics are likely to stay in school.

In order to analyze the effect of the exogenous variable on the probability of occurrence of the three alternative outcomes, we calculate marginal effects using all the models presented in Table 4. Results are found in Table 5.

We concentrate in the last three columns, which correspond to the model estimated with municipality dummies, to discuss the quantitative results. Since coffee production and the adult wage variables are measured in natural logarithms, we can interpret the marginal effects roughly as the impact of a 100 percent increase in the independent variables on the probability that each outcome occurs. For example, a 100 percent increase in the value of coffee production leads to an increase of 0.05 percentage points in the proportion of children working only (Column 1), to an increase of 0.15 percentage points in the proportion of children working and going to school (Column 2), and to a reduction of 0.20 percentage points in the proportion of children only going to school (Column 3). To get a sense of the magnitude of these changes, we divide the marginal effect by the observed probability of each outcome, which is reported at the bottom of Table 5. The estimated changes correspond to increases of 1.2 percent in the probability of work only and 0.9 percent in the probability of work and schooling, and to a reduction of 0.25 percent in the probability of schooling only.

In relation to the wage of the head of the household, a 100 percent increase is associate with a reduction of 0.09 percentage point in the fraction of children only working, a reduction of 2.9 percentage points in the fraction of children working and going to school, and an increase of 3 percentage points in the fraction of children only going to school. In terms of relative sizes of the groups, these magnitudes represent the following proportional changes: a reduction of 2.2 percent in the probability of work only, a reduction of 17.2 percent in the probability of work and schooling, and an increase of 3.8 percent in the probability of schooling only.

Overall, the responses of households to changes in economic activity and wealth follow the pattern predicted by theory. The shocks to local economic activity, once we control for household characteristics, socioeconomic status, time trends, and location-specific factors, do seem to isolate mostly the substitution effect from the increased demand for child labor. On the other hand, socioeconomic status and the wage of the head of the household seem to isolate the income effect, leading to better outcomes in terms of the allocation of time of children. The conflicting results found in previous papers are obtained here as different dimensions of the response of families to increases in income and in the level of economic activity.

B. Robustness

There are various potentially important dimensions of the problem that are not addressed explicitly in the model from Section II. In this subsection, we address some of these issues and also check the robustness of the results to alternative definitions of variables and sample. Table 6 presents the results, which follow the same specification from the last two columns of Table 4 (including municipality dummies and allowing coefficients to vary across equations). These results refer to estimations that change the definition of work, restrict the initial sample, or include further controls in our baseline specification.

Table 5
Household Choices of Child Labor and Schooling—Generalized Ordered Logits Marginal Effects (Specifications from Table 4), Brazil, 1993–2003

Coefficient	Model 1			Model 2		
	Work Only	Work and School	School Only	Work Only	Work and School	School Only
Coffee value (ln)	0.000517***	0.00515***	−0.00567***	0.000911***	0.00419**	−0.00510***
	[0.000168]	[0.00163]	[0.00179]	[0.000194]	[0.00187]	[0.00193]
Household wage (ln)	−0.00342***	−0.0341***	0.0375***	−0.00336***	−0.0343***	0.0376***
	[0.000398]	[0.00385]	[0.00418]	[0.000690]	[0.00428]	[0.00438]
Demographic characteristics		X			X	
Year fixed effects		X			X	
State fixed effects		X			X	
Municipal fixed effects						
Observations frequency	0.043	0.169	0.788	0.043	0.169	0.788

	Model 3			Model 4		
Coefficient	Work Only	Work and School	School Only	Work Only	Work and School	School Only
Coffee value (ln)	0.000648***	0.00445**	−0.00510***	0.000518***	0.00145	−0.00197
	[0.000150]	[0.00187]	[0.00192]	[0.000197]	[0.00281]	[0.00292]
Household wage (ln)	−0.00194***	−0.0361***	0.0380***	−0.000938**	−0.0291***	0.0300***
	[0.000556]	[0.00426]	[0.00435]	[0.000413]	[0.00429]	[0.00442]
Demographic characteristics		X			X	
Household wealth and infrastructure		X			X	
Year fixed effects		X			X	
State fixed effects		X				
Municipal fixed effects					X	
Observations frequency	0.042	0.0169	0.789	0.042	0.0169	0.789

Notes: Robust standard errors are in parentheses. *** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$. Unit of observation is child. Effect of change in independent variable on probability of outcome (specifications from Table 4). Independent variables are described in Table 4. Sample described in Table 4. Data are from PNAD and IBGE agricultural surveys.

Table 6
Dealing with Issues not Addressed in the Model—Household Choices of Child Labor and Schooling—Generalized Ordered Logits Marginal Effects, Brazil, 1993–2003

Coefficient	Work Defined as ≥ 15 Hours per Week		Sample of Sons and Daughters		Family Structure Controls	
	Work/Work and School	Work and School/School	Work/Work and School	Work and School/School	Work/Work and School	Work and School/School
Coffee value (ln)	−0.102***	−0.0243	−0.0947***	−0.0198	−0.0987***	−0.0142
	[0.0340]	[0.0231]	[0.0351]	[0.0240]	[0.0333]	[0.0238]
Household wages (ln)	0.117	0.207***	0.142*	0.243***	0.131*	0.244***
	[0.0751]	[0.0382]	[0.0770]	[0.0373]	[0.0725]	[0.0378]
Number of observations	26,098		25,468		26,293	

Coefficient	State-Specific Time Dummies		Municipality GDP p.c. and Population Controls		Top 40 Percent Coffee-producing Municipalities	
	Work/Work and School	Work and School/School	Work/Work and School	Work and School/School	Work/Work and School	Work and School/School
Coffee value (ln)	−0.107**	−0.0392*	−0.363***	−0.0455	−0.0604*	−0.0171
	[0.0429]	[0.0225]	[0.123]	[0.0405]	[0.0338]	[0.0333]
Household wages (ln)	0.154**	0.247***	0.105	0.288***	0.153*	0.245***
	[0.0753]	[0.0365]	[0.155]	[0.0562]	[0.0817]	[0.0400]
Number of observations	25,673		11,336		20,673	

Notes: Robust standard errors are clustered at municipality level in parentheses. *** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$. Unit of observations is child. Dependent variable is categorical, indicating whether child works only (0), works and goes to school (1), or goes to school only (2). These are the same specification as the last two columns of Table 4. Independent variables include (not shown): demographic characteristics (education, female dummy, race dummies, age, rural dummy, age of head of household, education of head of household, and female head of household dummy), household wealth and infrastructure (number of bedrooms per capita, dummies indicating whether household has electricity, phone, television, refrigerator, and washing machine, as well as tenure of head of household in current job, and ln of other household income), year dummies, and municipality dummies. "More than 15 hours of work" per week redefine children as working only if they worked at least 15 hours. "Sons and daughters" restricts sample to sons and daughters of the head of household or of her/his partner. Family structure controls include dummies for both parents working, single parent household, presence of siblings aged between 0 and 5, between 6 and 9, and between 15 and 18, number of siblings, and dummy for presence of elderly person. "Municipal GDP per capita and population" includes ln(gdp pc) and ln(population) as control. "Top 40 percent Coffee-producing Municipalities" restricts the sample to the top 40 percent municipalities among those with positive coffee production. State-specific time dummies include different time dummies for each state. Unless otherwise noted, the sample is restricted to children aged between 10 and 14 related to head of the household (heads of household are restricted to fully employed and aged between 18 and 65), and to municipalities with fewer than 100,000 inhabitants in the top 60 percent of coffee-producing municipalities. Data are from PNAD and IBGE agricultural surveys.

First, we change the definition of child labor to at least 15 hours of weekly work at some point during the previous year, in order to characterize a stronger attachment to the labor market. Our previous definition might be seen as including children without real attachment to the labor market, but that could be partly driving the results. Once we reestimate the model with this new definition of work, there is very little change. In fact, the majority of working children in our sample worked more than 15 hours a week at some point during the previous year (see Figure 1). Results are very similar to the ones from the last two columns from Table 4, but for the fact that the estimated coefficients for the coffee production variable tend to be slightly higher, and the ones on the wage of the head of the household slightly lower.

Following, we restrict the sample to only sons and daughters of the head of the household (or of her/his partner), to avoid the comparison of children who may be treated differently by the main family unit in the household (in the case of extended families living together). Results are virtually identical to those from Table 4.

In sequence, we include as additional controls (in both x_i and z_i) variables related to family structure, which are not modeled explicitly in our theory, but which may be important in determining allocation of time and resources within the household. These may be related to demand for household work, substitution of domestic or market labor across different members of the household or dilution of family resources across children. In order to account for these possibilities, we include as additional independent variables: a dummy indicating whether both parents work; a dummy indicating whether both parents live in the household; three dummy variables indicating the presence of siblings aged between 0 and 5, 6 and 9, and 15 and 18; a dummy indicating the presence of and elderly person in the household; and a variable indicating the total number of children in the household. The main results are again very similar to those presented before. Among the family structure variables, number of children in the household, the presence of a sibling aged between 15 and 18, and the fact that both parents work are all associated with worse allocations of time for children, toward more work and less schooling. The other variables do not have robust and significant effects. So, despite the fact that family structure indeed seems important in determining the allocation of children's time, it does not affect our previous results in any systematic way.

The next four columns try to account for state or municipality factors that may be simultaneously correlated with the value of coffee production and with child labor. First we include state-specific time dummies, in order to account for the possibility of differential trends in coffee production across the different areas of the country, maybe reflecting differential development trends also associated with child labor. Quantitative results remain very similar to the other columns from Table 6 and, if anything, significance and magnitude of some coefficients becomes stronger.

Second, in order to address the possibility that the value of coffee production reflects in part different municipality characteristics, we control for municipality GDP per capita and population (natural logarithm). Local GDP per capita may indicate the dynamism of the local economy and, therefore, may be correlated with both coffee production and child labor. The series for municipality level GDP per capita for Brazil is much shorter than the other datasets used in the paper and, therefore, our sample is greatly reduced when this variable is included. Interestingly, once we control for GDP per capita, the pattern initially observed in the first columns

of Table 6 becomes even more evident: The coefficient on the value of coffee production becomes much larger for families choosing between *work only/work and school*, when compared to families choosing between *work and school/school only*.

In the final two columns of Table 6, we change the sample to include only the top 40 percent among coffee-producing municipalities. Again, the sample is reduced substantially. Despite some reduction in magnitude on the estimated effect of coffee production, qualitative results remain again very similar.

As a final exercise in this subsection, we estimate our baseline specification from the last two columns of Table 4 for boys and girls separately, and also for children in an older age group (15 to 18, as compared to 10 to 14 in our original sample). Results are presented in Table 7. As one might expect, in the younger age group boys' responses to shocks to local economic activity tend to be more elastic than girls'. In relation to the response of the different genders to changes in the wage of the head of the household, there is no clear pattern. Girls seem to be more sensitive to the wage of the head of the household at younger ages, but results are reversed in the age group between 15 and 18, where differences are relatively modest. Overall, older kids do not appear to be responsive to shocks to local economic activity, maybe because they already have a stronger attachment to the labor market. In any case, quantitative and qualitative results across genders for children between 10 and 14 are similar to the evidence presented before.

Going back to Table 6, the results seem to confirm the pattern that already appeared in Table 4, suggesting a stronger impact of the value of coffee production on poorer families, who are facing the margin of choice *work only/work and school*.

C. Instrumental Variables

We conduct one final exercise, adopting yet an alternative approach to deal with the source of variation in the value of coffee production. We want to further explore this issue here in order to have enough confidence that our coffee production variable is indeed capturing exogenous shocks to the local demand for low-skill labor. This alternative strategy uses climatic variables to instrument for the value of coffee production. We have yearly data at the municipality level on average temperatures and rainfall by quarters (December to February, March to May, June to August, and September to November). The dimension of coffee production explained by changes in rainfall and temperature is likely to be driven by temporary climatic shocks, and less subject to the problem of being correlated with intrinsic and time invariant characteristics of municipalities.

There are limitations to the use of instrumental variables in nonlinear models that are beyond the scope of this paper (see, for example, Angrist and Pischke 2009). We do not deal with these issues here. We simply estimate a linear first stage and, analogously to 2SLS, plug-in the fitted values from the first stage in a nonlinear second stage, taking the precaution of estimating standard errors for the second stage via bootstrapping of the entire process (the exact procedure is detailed below). Despite its lack of rigor, we still believe that this strategy sheds light on the driving forces behind the results presented before. In addition, the use of climatic variables as instruments for agricultural activity is, on itself, something novel in the child labor literature.

Table 7
Results by Gender and Age—Household Choices of Child Labor and Schooling—Generalized Ordered Logits, Brazil, 1993–2003

Coefficient	Girls 10–14		Boys 10–14	
	Work/Work and School	Work and School/School	Work/Work and School	Work and School/School
Coffee value (ln)	−0.0868**	0.015	−0.107**	−0.0337
	[0.0365]	[0.0293]	[0.0522]	[0.0255]
Household wage (ln)	0.221*	0.307***	0.152*	0.228***
	[0.128]	[0.0560]	[0.0853]	[0.0430]
Number of observations	12,759		13,534	

Coefficient	Girls 15–18		Boys 15–18	
	Work/Work and School	Work and School/School	Work/Work and School	Work and School/School
Coffee value (ln)	0.0083	−0.0148	0.0105	−0.0335
	[0.0350]	[0.0233]	[0.0179]	[0.0222]
Household wage (ln)	0.139**	0.324***	0.123***	0.376***
	[0.0556]	[0.0448]	[0.0412]	[0.0362]
Number of observations	7,617		9,588	

Notes: Robust standard errors clustered at municipality level in parentheses. *** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$. Unit of observation is child. Dependent variable is categorical, indicating whether child works only (0), works and goes to school (1), or goes to school only (2). Independent variables include (not shown): demographic characteristics (education, female dummy, age, rural dummy, age of head of household, education of head of household, and female head of household dummy), household wealth and infrastructure (number of bedrooms per capita, dummies indicating whether household has electricity, phone, television, refrigerator, and washing machine, tenure of head of household in current job, and ln of other household income), year dummies, and municipality dummies. Data from PNAD and IBGE agricultural surveys. Sample restricted to children related to head of the household (heads of household restricted to full employed and aged between 18 and 65), to municipalities with less than 100,000 inhabitants in top 60 percent of coffee-producing municipalities, and to boys or girls, and ages 10 to 14 or 15 to 18, according to what is indicated in the table.

We run a first-stage regression where the dependent variable is the value of coffee production per capita at the municipality level, the instruments are the climatic variables described above, and additional controls include all the right-hand side variables used in the final specification from Table 4. The value of coffee production and the climatic variables are observed at the municipality level, but since our second stage uses individual level data, we run the first stage also at the individual level, allowing for clustering of standard errors by municipality. Table 8 presents the results from this first-stage estimation. The first column displays the coefficients on the climatic variables when they are all included together (coefficients on other independent variables are omitted). Three of the climatic variables turn out nonsignificant and the joint F-statistic on the instruments is 9.58. In order to avoid a potential problem of weak instruments, we rerun this first-stage regression including only the climatic variables that appeared as significant. This result is presented in the second column of Table 8. The joint F-statistic on the instruments in this case is 14.93. This is the first stage we use in our instrumental variables estimation. Standard errors are calculated by bootstrapping the entire process (first and second stages together, 500 repetitions), using municipalities as resampling clusters.

The results from the instrumental variables estimation are presented in Table 9. Results are quite similar to those obtained before. The pattern present in Table 9 is of a much stronger effect of the coffee production shock on the decision margin *work only/work and school*. This suggests that poorer families, where the children are usually engaged in work, are the ones most affected by the increase in the demand for low-skill labor represented by the shock to coffee production. This coefficient is negative and significant, and larger in absolute value than most of those presented in Table 6. The results for the margin of choice *work and school/school only* are also negative, but not statistically significant. The wage of the head of the household, on its turn, appears as positive and significant in both transition equations, being larger in the transition *work and school/school only* than in the transition *work only/work and school*, also similarly to what was obtained before.

Overall, results are remarkably consistent. Shocks to local economic activity captured by the exogenous component of the value of coffee production do seem to increase the opportunity cost of children's time, worsening their allocation of time (moving from schooling toward more labor). This effect is particularly strong for poorer households with working children who are considering whether or not send the children to school. Still, pure income effects on child labor, as captured by increases in the wage of the head of the household, are associated with improvements in the allocation of children's time, toward more schooling and less labor. This pattern of results is clearly present even when we use municipality fixed effects or instrument the value of coffee production with climatic variables.

V. Concluding Remarks

In this paper, we develop a simple theory of household choices of child labor and schooling. The theory characterizes the household problem in a way that can be represented by a generalized ordered discrete choice model. In this

Table 8

First Stage—Climatic Determinants of Coffee Production—Individual Level
Regressions, OLS, Brazil, 1993–2003

Dependent Variable: Value of Coffee Production per capita (ln)

Temperature Dec-Feb	−0.529**	−0.420**
	[0.241]	[0.205]
Temperature Mar-May	0.181	
	[0.229]	
Temperature Jun-Aug	−0.670***	−0.590***
	[0.176]	[0.174]
Temperature Sep-Nov	0.414**	0.418***
	[0.168]	[0.153]
Rain Dec-Feb	−0.00110	
	[0.00354]	
Rain Mar-May	0.00705**	0.00691**
	[0.00280]	[0.00311]
Rain Jun-Aug	−0.0115**	−0.0119*
	[0.00575]	[0.00601]
Rain Sep-Nov	0.00146	
	[0.00409]	
F of instruments	$F(8,142) = 9.58$	$F(5,142) = 14.93$
	p-value $= 0.00$	p-value $= 0.00$
Number of observations	24705	24705
R Squared	0.309	0.308
Number of municipalities	143	143

Notes: Robust standard errors clustered at the municipality level in parentheses. *** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$. Unit of observation is child. Dependent variables are value of municipality coffee production per capita (ln). Instruments are yearly observations on average temperatures and rainfall by quarters (Dec-Feb, Mar-May, Jun-Aug, and Sep-Nov). Other independent variables include all controls from Table 4. Sample restricted to children aged between 10 and 14 related to household (heads of household restricted to full employed and aged between 18 and 65), and to municipalities with less than 100,000 inhabitants in top 60 percent of coffee-producing municipalities. Data from PNAD, IBGE agricultural surveys, and from the Climate Research Unit of the University of East Anglia.

model, families can choose among three alternatives for children: work only, work together with schooling, and schooling only.

Given the limitations of the data typically available, we introduce some simplifying assumptions and apply an empirical specification inspired by the theoretical model to Brazilian data. This specification uses shocks to local coffee production as a way to distinguish between the roles of increases in family wealth (income effect) and in the opportunity cost of children's time (substitution effect) in determining the

Table 9

Instrumental Variables Estimates—Household Choices of Child Labor and Schooling—Generalized Ordered Logits, Brazil, 1993–2003

Coefficient	Instrumental Variables	
	Work/Work and School	Work and School/School
Coffee value (ln)	−0.1459***	−0.0462
	[0.0554]	[0.0542]
Household wage (ln)	0.2254***	0.2961***
	[0.0653]	[0.0374]
Number of observations	23,634	

Notes: Robust standard errors clustered at municipality level in parentheses. *** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$. Unit of observation is child. Dependent variable is categorical, indicating whether child works only (0), works and goes to school (1), or goes to school only (2). Independent variables include all controls from Table 4 (not shown). Instruments are quarterly average temperatures and rainfall (first stage is Column 2 in Table 8). Standard errors for the IV estimation calculated through bootstrapping of the entire two-stage IV process (500 replications), accounting for clustering of observations at the municipality level. Sample restricted to children aged between 10 and 14 related to head of the household (heads of household restricted to full employed and aged between 18 and 65), and to municipalities with less than 100,000 inhabitants in top 60 percent of coffee-producing municipalities. Data are from PNAD, IBGE agricultural surveys, and from Climate Research Unit of the University of East Anglia.

incidence of child labor. We find that household characteristics associated with higher permanent income and wealth (or with less dependence on child's income) are associated with lower incidence of child labor and higher school attendance. At the same time, conditional on household wealth and socioeconomic characteristics and on long-term trends, increases in labor demand due to shocks to local economic activity increase the opportunity cost of children's time, therefore increasing the incidence of child labor and reducing school attendance. These results hold even with municipality fixed effects and when the value of coffee production is instrumented with climatic variables.

Our approach trusts more heavily on theory than previous empirical work on child labor. As a consequence, we are able to understand the reasons behind some of its seemingly conflicting results. Our evidence related to family wealth and increases in the demand for child labor seems to isolate, respectively, the income and substitution effects present in different types of income variation. The paper shows that, in order to fully understand the consequences of a certain change in income or in the level of economic activity, one must understand how it affects the full income of households and the opportunity cost of children's time.

Appendix

Table A1

Simple Logits and Municipality Level OLS for Work and Schooling Separately, Brazil, 1993–2003

Coefficient	Individual Level Data		Municipal Level Data	
	Work	School	Work	School
Coffee value (ln)	0.0178	− 0.0181	− 0.000134	− 0.000918
	[0.0232]	[0.0204]	[0.00358]	[0.00110]
Household wage (ln)	− 0.168***	0.111**	− 0.000583	− 0.00218
	[0.0361]	[0.0556]	[0.0263]	[0.0111]
Number of observations	27,413	27,413	1,268	1,268
R squared			0.559	0.491

Notes: Robust standard errors clustered at municipality level in parentheses. *** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$. Unit of observation is child in the first two columns and municipality in the last two. Dependent variable is either categorical indicating whether child works or goes to school Columns 1 and 2, respectively), or proportion of children working or going to school (Columns 3 and 4, respectively). Independent variables include (not shown): demographic characteristics (education, female dummy, race dummies, age, rural dummy, age of head of household, education of head of household, and female head of household dummy), household wealth and infrastructure (number of bedrooms per capita, dummies indicating whether household has electricity, phone, television, refrigerator, washing machine, as well as tenure of head of household in current job and ln of other household income), year dummies, and municipality dummies. The sample is restricted to children aged between 10 and 14 related to head of the household (heads of household restricted to full employed and aged between 18 and 65), and to municipalities with less than 100,000 inhabitants in top 60 percent of coffee-producing municipalities. Data are from PNAD and IBGE agricultural surveys.

References

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics—An Empiricist's Companion*. Princeton: Princeton University Press.

Baland, Jean-Marie, and James A. Robinson. 2000. "Is Child Labor Inefficient?" *Journal of Political Economy* 108(4):663–79.

Barros, Ricardo, Rosane Mendonça, and Tatiana Velazco. 1994. "Is Poverty the Main Cause of Child Work in Urban Brazil?" Working Paper 351, Instituto de Pesquisa Econômica Aplicada. Rio de Janeiro: Instituto de Pesquisa Econômica Aplicada.

Basu, Kaushik, and Pham Hoang Van. 1998. "The Economics of Child Labor." *American Economic Review* 88(3):412–27.

Beegle, Kathleen, Rajeev H. Dehejia, and Roberta Gatti. 2006. "Child Labor and Agricultural Shocks." *Journal of Development Economics* 81(1):80–96.

Black, Dan A., Terra G. McKinnish, and Seth G. Sanders. 2003. "Does the Availability of High-Wage Jobs for Low-Skilled Men Affect Welfare Expenditures? Evidence from Shocks to the Steel and Coal Industries." *Journal of Public Economics* 87(9–10):1921–42.

Bourguignon, François, Francisco H. G. Ferreira, and Phillippe G. Leite. 2003. "Conditional Cash Transfers, Schooling, and Child Labor: Micro-Simulating Brazil's Bolsa Escola Program." *World Bank Economic Review* 17(2):229–54.

Cardoso, Eliana, and André P. Souza. 2004. "The Impact of Cash Transfers on Child Labor and School Attendance in Brazil." Working Paper 04-W07, Vanderbilt University. Nashville: Vanderbilt University.

Dessy, Sylvain E. 2000. "A Defense of Compulsive Measures against Child Labor." *Journal of Development Economics* 62(1):261–75.

Duryea, Suzanne, and Mary P. Arends-Kuenning. 2003. "School Attendance, Child Labor, and Local Labor Markets in Urban Brazil." *World Development* 31(7):1165–78.

Edmonds, Eric V. 2005. "Does Child Labor Decline with Improving Economic Status?" *Journal of Human Resources* 40(1):77–99.

———. 2006. "Child Labor and Schooling Responses to Anticipated Income in South Africa." *Journal of Development Economics* 81(2):386–414.

Emerson, Patrick M., and André P. Souza. 2003. "Is There a Child Labor Trap? Intergenerational Persistence of Child Labor in Brazil." *Economic Development and Cultural Change* 51(2):375–98.

Gunnarsson, Victoria, Peter F. Orazem, and Mario A. Sánchez. 2006. "Child Labor and School Achievement in Latin America." *World Bank Economic Review* 20(1):31–54.

Kruger, Diana. 2006. "Child Labor and Schooling During a Coffee Sector Boom: Nicaragua 1993–1998." In *Trabajo Infantil: Teoría y Evidencia desde Latinoamerica*, ed. Luis F. Lopez-Calva, 319–343. Mexico City: Fondo de Cultura Económica de México.

———. 2007. "Coffee Production Effects on Child Labor and Schooling in Rural Brazil." *Journal of Development Economics* 82(2):448–463.

Levison, Deborah, Karine S. Moe, and Felicia Marie Knaul. 2000. "Youth Education and Work in Mexico." *World Development* 29(1):167–88.

Neri, Marcelo C., and Mark Thomas. 2001. "Household Responses to Labor-Market Shocks in Brazil, 1982–99." Unpublished Manuscript, Fundação Getúlio Vargas. Rio de Janeiro: Fundação Getúlio Vargas.

Ranjan, Priya. 2001. "Credit Constraints and the Phenomenon of Child Labor." *Journal of Development Economics* 64(1):81–102.

Rogers, Carol Ann, and Kenneth A. Swinnerton. 2004. "Does Child Labor Decrease When Parental Income Rise?" *Journal of Political Economy* 112(4):939–46.

Schultz, T. Paul. 1985. "Changing World Prices, Women's Wages, and the Fertility Transition: Sweden, 1860–1910." *Journal of Political Economy* 93(6):1126–54.

Williams, Robert. 2006. "Generalized Ordered Logit/Partial Proportional Odds Models for Ordered Dependent Variables." *Stata Journal* 6(1):58–82.