

Child Labor

Christopher Udry

Child labor is an insidious evil. Leaving aside pathological cases of child abuse and abandonment, it exists because it is the best response people can come up with to intolerable circumstances. It is particularly dangerous because it involves the sacrifice of a child's future welfare in exchange for immediate benefit, and difficult to combat because it involves questions of agency and power within households.

Some forms of child labor amount to direct abuse. There are children who work in dangerous conditions, in the sex industry, or in forced labor. But the vast majority of working children are engaged in less extreme activities, often on their own family's farm or business. For these children, the primary cost of child labor is the associated reduction in investment in their human capital. This occurs chiefly because child labor interferes with schooling. Not all work by children has this effect; I exclude such work from consideration in this essay and define child labor as the sacrifice of the future welfare of the child in exchange for additional current income. Although there are important challenges associated with empirically distinguishing child labor from the unproblematic light work that is an important component of rearing a child, we will see below that it is possible to design programs that specifically target child labor.

Many economists argue that child labor is a symptom of poverty and that its reduction can most effectively be accomplished through the alleviation of poverty. It is surely correct that child labor is a symptom of poverty; rarely do well-off parents sacrifice their children's education by sending them to work. However, child labor is also a cause of future poverty, so direct mea-

asures to move children from work into school can make an important contribution to poverty alleviation and to development in general.

In order to construct effective policies to address the problem of child labor, it is necessary to understand the circumstances that lead parents to send their children to work. That is the purpose of this essay. I make no attempt to survey the economic literature on child labor; Basu (1999) already provides an excellent review.

In the second section, I briefly describe some of the main features of child labor in developing countries. Poverty and child labor are mutually reinforcing; because their parents are poor, children must work and therefore remain out of school. As a consequence, these children grow up to be poor as adults, and the cycle continues. In the third section, I discuss the first of two features of child labor that give it a central place in a vicious cycle of poverty. This is the fact that the primary costs of child labor are realized so far in the future. When financial markets are poorly developed, the separation in time between the immediate benefits and long-delayed costs of sending children to work can result in too much child labor. The second feature is that the costs and benefits of child labor are not only separated in time; they are borne by different people: the child suffers the main consequences, while other household members benefit. This problem of *agency* is discussed in the fourth section. Finally, the fifth section concludes with a discussion of child labor policies.

PATTERNS OF CHILD LABOR

The International Labour Office (2002) estimates that about 210 million children between the ages of five and fourteen were working in 2000, about half of them working full-time. That implies that approximately one in ten of the world's children were working full-time. At the same time, UNESCO estimates that about one of every five primary school-aged children were not enrolled in school. The absolute numbers of children working are largest in Asia, but the incidence of child labor seems to be highest in Africa; the ILO estimates that about one third of children are economically active in Africa, about one sixth working full-time.

Child labor is overwhelmingly a rural and agricultural phenomenon. For example, in Pakistan, 70% of working children are employed in agriculture (Pakistan FBS 1996). Boys are more likely to work than girls, and older children are much more likely to be employed than their younger siblings (Grootaert and Patrinos 1999).

Our concern focuses on child labor that involves the sacrifice of a child's future welfare in exchange for a current benefit to the household. This is clearly the notion that motivates most child labor policy, and lies behind ILO Convention number 138.¹ The benefits to the household of sending a child to work are the wages of that child (or, equivalently, the increased production on the family farm) and the reduced education expenditures

from not sending her to school. The primary costs of child labor are the lower future earnings of the child when she enters the adult labor market with lower educational attainment. In addition, there is very strong evidence of important nonmarket returns to education in child rearing: the children of well-educated parents, particularly mothers, are healthier. The sacrifice of these returns should also be counted as a cost of child labor. Finally, there are benefits to education, and thus costs to child labor, that extend beyond the immediate family. Educated coworkers may improve the productivity of everyone, and a well-educated populace provides a vital foundation for a vibrant society. These benefits that accrue to people outside of the household should also be counted as part of the cost of child labor, but I will for the most part ignore them in this essay. The primary thrust of my argument is that there are reasons to expect child labor to be “too high” among poor families in developing countries, even apart from the benefits of education for the broader public. Taking these into account would only strengthen the argument.

There are certain well-established empirical regularities about child labor that should inform our discussion. First, it is clear that child labor overwhelmingly reflects the poverty of the households in which the children live. Fallon and Tzannatos (1998) review a variety of studies that indicate a strongly negative relationship between the incidence of child labor and household income, but this relationship is less marked in more affluent developing countries. Ray (2000) finds a strong negative correlation between household income and child labor, and a positive relationship between household income and school enrollment in Pakistan, but no such relationship in relatively wealthier Peru.

It is important to note that the strong empirical evidence that child labor declines and school enrollment increases with household income does *not* imply that increases in wages necessarily reduce child labor. When a household becomes better off, it tends to increase the school enrollment of its children. Economists call this a positive “income effect.” The children in a household made better off through, for example, an unconditional government grant will tend to work less and attend school more. However, when the additional income comes from an increase in wages, a countervailing force also exists. Wages of adults and children tend to move together, and an increase in child wages increases the effective cost to the household of sending a child to school rather than to work. Each hour the child spends in school (and thus not working) reduces the household’s current consumption by more when the child’s wage is higher. This “substitution effect,” therefore, tends to increase the incidence of child labor. If this negative substitution effect is sufficiently strong, it could outweigh the positive effect of the household’s increased income on child schooling. Whether it does or not depends on circumstances. For example, Kruger (2002) shows that child labor increases and school attendance decreases as coffee prices—and thus the returns to child labor—increase in Brazil. In contrast, Edmonds and Pavcnik

(2002) show that in Vietnam, increases in rice prices were strongly associated with declines in child labor.

It is also clear that child labor has important detrimental effects on schooling attainment, and thus on the future income of children. As already noted, not all work by children has this effect. Ideally, such benign work by children (occasional light work on the family farm, or limited household work) is excluded from data collection on child labor. An important question to resolve is the extent of work by children that does interfere with schooling and thus future earnings. How many of the ILO's 210 million working children are sacrificing their education? This is inherently a difficult question to answer, because child labor and school enrollment influence one another. Is a child not in school because she is working? Or is she working because there are no good opportunities for her to be schooled? Or is she sacrificing leisure to both work and attend school? However, the existing evidence is strong. For example, Psacharopoulos (1997) shows much lower educational attainment by children who work in Venezuela and Boliva. Using a very different methodology, Boozer and Suri (2001) find similar results for child labor and school attendance in Ghana.

Households that are very poor are much more likely to send their children to work, and child labor contributes to poverty in the next generation by reducing schooling attainment. This circular pattern of positive feedback between poverty and child labor may lead to a vicious cycle of poverty, in which the descendants of the poor remain poor because they were poorly educated. This cycle can be the foundation of a classical "poverty trap." However, if the cycle can be broken, the same positive circular causation can contribute to a takeoff into sustained growth. If schooling attainments can be improved, then the next generation's income is higher and their children can in turn become yet better educated. It is essential, therefore, to understand the specific mechanisms that can trap people in the awful equilibrium of persistent poverty, excessive child labor, and low education over generations.

The crucial mechanisms are, first, an inability to seize advantageous long-run investments in children's human capital because of credit market constraints and, second, problems of agency within households. These two mechanisms operate simultaneously and can interact in important ways.

IMPERFECT FINANCIAL MARKETS, CHILD LABOR, AND INVESTMENT IN HUMAN CAPITAL

I begin by abstracting from any problems associated with agency, and assume that parents fully internalize the costs of sending their children to work. From the point of view of society, what is the appropriate level of child labor? Suppose, to start, that the costs of working and the benefits of schooling are entirely private; that is, they are limited to the increased productivity, and therefore income, of better-educated adults. For now, therefore, we are ig-

noring the benefits of well-educated individuals for the rest of society. The benefit of child labor is the current wage earned by the child (and the reduced cost of schooling). These benefits are realized immediately, so no discounting is required. The costs of additional child labor are the lower wages that the child receives when she grows up less well-educated because she worked as a child. These costs are realized in the far future, so for cost-benefit analysis we calculate the present discounted value of these costs. For a given absolute cost, a higher interest rate implies a lower present discounted value of the cost.

From a social point of view, it is efficient to increase child labor and reduce schooling up to the point at which the present discounted value of future costs of additional child labor are just balanced by the current benefit to the household of that additional child labor. It need not be the case that the socially efficient level of child labor is zero; this will depend upon the productivity of child labor, the degree to which schooling improves future productivity, and the interest rate at which future earnings are discounted. For example, if a child already has sufficient schooling so that further years of education have a relatively small impact on her future income, and if she could generate a lot of income by working, and if interest rates are relatively high, then the immediate benefit of having the child work might be sufficiently large to offset the present discounted value of her future lower earnings as a less well-educated adult. On the other hand, if further schooling for the child would greatly increase her future income, and if the current income she could earn by working is relatively low, and if interest rates are low, then the immediate benefit of the child working would not be sufficient to outweigh the present discounted value of her future higher income from attending school, and from a social point of view she should stay in school.

If financial markets operate smoothly and there are no issues of agency, this is precisely the calculus that will guide the decisions of parents regarding work and school for their children. In this case, even if parents are poor, perfect credit markets permit them to borrow to finance the education of their children, confident in their ability to repay the loan out of the increased earnings of their well-educated adult children. These private decisions will be socially optimal.

Obviously, if the rest of society benefits from children's improved education, then these benefits will not be fully taken into account by parents as they invest in their children's schooling. In this case, from a social point of view the level of child labor will be too high, and the level of schooling too low even if financial markets operate smoothly. The existence of these social benefits from schooling is an important element in the traditional argument for subsidization of education and public schooling.

However, even if these social benefits are unimportant, in the real world financial markets are not sufficiently well developed to support the optimal calculus described in the first three paragraphs of this section. A parent who is unable to smoothly transfer income from the far future into the present

by borrowing will choose too high a level of child labor. Consider a very simple example, in which the parent can borrow, but only at an interest rate that is higher than is relevant for social decision-making. This would occur, for example, if the parent can borrow only from a monopolistic money-lender. At this higher interest rate, the present discounted value to the household of the future costs of child labor are lower than they are to society as a whole, and so the child works more and attains a lower level of schooling.

This reasoning holds *a fortiori* when the parents have no or only constrained access to credit. In this case the parents trade off the (far) future costs of child labor against the immediate benefit of child labor to the household, without the possibility of easing that trade-off by transferring resources from the future. If the household is too poor, the value of the immediate return from the child's working trumps the future higher wages the child could earn if she stayed in school, and the child is sent to work.²

There is little doubt that inadequate access to financial markets is a barrier to investment in education and a force pushing the children of many poor households into the labor market. In conversation parents testify to their importance. To this, economists have added quantitative evidence. However, to be persuasive, they must go beyond the well-established correlation between poverty and child labor. To see why, consider two families who are similar, except that one is poorer than the other. In particular, each has a twelve-year-old boy who could earn \$1/day in a local business, or who could go to school. Suppose we see that the child of the less poor family is in school, while the child of the poorer family is working. This is certainly consistent with the argument that the poorer family has inadequate access to financial markets, and faces a higher interest rate than does the better-off family. At that higher interest rate, the present discounted value of the future gain to income by sending the child to school might not offset the \$1/day he can earn now. On the other hand, there could be many other differences between these two families. For example, the better-off family might have social connections that increase the value of schooling to them. Or the poorer family might have access only to a lower-quality school. Broadening our vision beyond the narrow focus on monetary costs and benefits, one family might be better off because the parents are better-educated themselves, and their education leads them to place a higher value on schooling. The point is that there are many reasons why we might see a correlation between low income and a high incidence of child labor. Navigating the causal pathways is difficult.

This is a general problem in the social sciences. The strategies that economists have used to attempt to distinguish various causal explanations for this particular correlation provide useful examples of the general principle.³ In this case, some of the more persuasive evidence comes from a series of studies that relate both school attendance and child labor to transitory income shocks afflicting households that can be attributed to random events outside the control of the household. The importance of the last two phrases

is that these studies aim to identify households that are temporarily poor (or rich) for no particular reason—it was just an act of God. If this can be done, we can argue that the *only* difference between households that have a temporary good shock and those that have a temporary bad shock is that (on average) those with a good shock are relatively well off.

To understand the importance of such evidence, consider a hypothetical household engaged in farming in a developing country. Like all farmers, this family is subject to random, transitory shocks that affect its income (think of shocks such as a localized flood). An important consequence of well-developed financial markets is that this family's decisions regarding the education and labor force participation of its children would be entirely unaffected by the realization of such transitory shocks. The present discounted value of the future costs of child labor are unaffected by any temporary production shock. The immediate benefits are unaffected as well, as long as the shock is sufficiently localized that the wage for child labor does not change.⁴ Therefore, if this family has access to smoothly operating credit markets, it will simply borrow (or draw down its savings) to maintain its base level of consumption despite the adverse shock, and the schooling and labor status of its children would remain unchanged.

This conclusion no longer holds when the family cannot borrow and does not hold savings over the long term. Now, faced with an adverse transitory shock, a poor household is forced into a stark choice: maintain the school enrollment of the children and face a decline in an already inadequate level of consumption, or try to protect the family's current living standard by relying on increased child labor. Moreover, if we are convinced that these adverse shocks are distributed randomly through the population, and that the only effect of these shocks is to temporarily lower income, then we can be confident that the causal connection truly does run from income to schooling and child labor.

Beegle et al. (2002) provides a good example of this approach. They find that rural Tanzanian children work more when their family farms have experienced adverse events such as fires or loss due to insects or rodents. The survey they use followed households over three years, so they can also show that this effect is temporary: children work more when their household is poorer as a consequence of suffering from these events, then they work less when the household recovers. Finally, they show that the sensitivity of child labor to these shocks is greatest for poorer households with fewer assets that can serve as collateral, thus providing further support for the hypothesis that imperfect financial markets play a central role in determining the amount of child labor in rural Tanzania. In a complementary paper, Jacoby and Skoufias (1997) find that school attendance drops among children in three Indian villages when their households' income temporarily drops. This is not simply a rural phenomenon, however. Duryea et al. (2003) show that urban children in Brazil are much more likely to enter the workforce when their father becomes unemployed. For example, the probability of a sixteen-year-old girl

in São Paulo entering the labor force in the next year jumps from 22% to 35% if her father loses his job; simultaneously, her probability of continuing in school drops from 70% to 60%.

Child labor can be seen as a draconian choice made by poor households faced with severely limited options. Even if the parents are fully altruistic toward their children, in the sense that they treat the future costs to the children of current child labor symmetrically with current benefits to the household, poorly functioning financial markets can induce too much child labor and too little schooling. The benefits of child labor are realized immediately. Without access to credit markets, poor households may find it too difficult to sacrifice these immediate benefits in order to reduce the far future costs associated with child labor. It bears emphasizing that this is not a consequence of impatience or an unwillingness on the part of poor households to plan for the future; rather, it is a reflection of poverty and inadequate access to capital markets.

AGENCY

Decisions regarding child labor and schooling are generally made by parents. This raises issues of agency, because decisions are being made by individuals who do not necessarily themselves experience the full implications of these decisions. Even if parents are altruistic toward their children—and surely this is the case for the vast majority of families—issues of bargaining and negotiation within households, and the difficulty of making commitments that bind over generations, may make it difficult to achieve optimally low levels of child labor.

First, consider a case in which agency causes no deviation from the socially efficient levels of child labor and schooling, in the spirit of the classic “rotten kid” theorem of Becker (1974). Suppose that the parent feels altruistic toward the child, in the sense that the parent’s welfare increases when the child’s welfare increases, and that the parent has access to perfect financial markets. In addition, suppose that the parent expects to leave a positive bequest to the child. In this case, the parent will choose to set the level of child labor to the socially optimum level, as described above. The argument is quite simple: the parent would like to help the child achieve a particular level of welfare, and the parent has two tools available to do so: the parent chooses the amount of child labor (and thus determines the level of schooling for the child), and the parent can give the child a bequest. The parent will choose the minimal-cost means of achieving any given level of child welfare; to do otherwise would waste resources that could be used to achieve higher welfare for the child, the parent, or both. If the parent chooses a level of child labor greater than is socially optimal, he will be wasting resources. He can reduce child labor a bit, reduce the future bequest left to the child to compensate, and have money left over to increase everyone’s welfare. Therefore, a parent who cares about the welfare of his child *and* who plans

to leave a positive bequest to that child would ensure that the child's labor force participation matches the socially efficient level.

However, suppose that the parent plans to leave no bequest. This is most likely to occur in a poor family, particularly in a family in which the parent's generation is especially poor relative to future generations. Child labor in this circumstance will be inefficiently high and schooling attainment too low, because once bequests have been reduced to zero, this is the only instrument available to the parents to transfer resources from the next generation to support current welfare (Baland and Robinson 2000). A potential way to reduce child labor would be for the parent to borrow to finance current consumption, with the child committing herself to pay back the loan from her future higher earnings. However, such intergenerational contracts are not enforceable.

Therefore, even when financial markets operate perfectly smoothly (at least within generations) and parents are altruistic toward their children, agency problems can induce too much child labor and too little investment in education. The source of the problem is that poor parents who plan to leave no bequest to their children use child labor to support the current consumption of the household.

Agency problems become even more salient when they occur in the typical environment of imperfect financial markets. If the household cannot borrow (and does not plan to save), then decisions regarding child labor and educational investments cannot be made by balancing the current financial gain and discounted future financial cost of child labor. Instead, decisions are made by balancing *subjective* welfare costs and benefits. Parents balance the benefit in terms of current welfare of increasing child labor (and reducing schooling) against the current subjective cost of the child's future reduced welfare.

The immediate question, of course, is *whose* subjective welfare determines the child's education and labor force participation? The two parents might hold divergent views about these costs and benefits. In fact, there is mounting evidence that this is so, and that these divergent opinions can have important effects on child welfare.

Until fairly recently, economists had ignored issues of agency within households, relying on what has come to be called the "unitary household model." This model assumes that the choices made in households can be treated as if they were made by a single individual. There was never much of a theoretical justification for this assumption; it was made for convenience, driven by the fact that data (particularly from official statistical agencies) tends to come in household-sized chunks.

An important implication of the unitary household model is that income is pooled. Whether extra income comes from the husband or the wife is irrelevant for decisions regarding expenditure or investment in children; it's all just extra income for the household. When this implication is examined using data, it is almost universally rejected. For example, Duflo (2003) finds

that the nutrition of girls in South Africa is dramatically improved when their grandmothers receive old-age pensions, but is entirely unaffected when the pension is received by their grandfathers. This and much additional evidence implies that the unitary household model is an inappropriate building block for thinking about decisions within the household regarding investment in their children. Parents may have divergent preferences regarding such investments, so that shifts in bargaining power within the household could have important effects on child labor.

Economists are far from a general understanding of intrahousehold bargaining processes. In fact, the dominant successor model to the unitary household model is deliberately agnostic regarding these negotiations, assuming only that the household efficiently uses all the resources available to it.⁵ However, some general patterns have emerged from a fairly lengthy sequence of empirical studies. In particular, researchers have found that extra income in the hands of mothers is associated with higher levels of investment in child human capital (see Haddad et al. 1997).

LESSONS FOR POLICIES THAT CAN MOVE CHILDREN FROM WORK TO SCHOOL

Child labor should be understood as the consequence of people coping with extreme circumstances. It is a result of current poverty, and a cause of continued poverty for the children who sacrifice their education in order to work. It is a particularly insidious problem because its primary costs are long-delayed and realized by the child, while the benefits are immediate and directly affect decision-makers within the household.

We know that the ultimate instrument for the elimination of excess child labor is the alleviation of poverty. The evidence is indisputable: child labor as a mass phenomenon disappears when the population moves out of poverty. While this is a sure solution, we're not willing to wait.

The obvious response is an outright ban of the practice of child labor. The first difficulty is that it is by no means clear that developing country governments have tools available to enforce such a ban. The task would be extremely difficult, because most child labor is in agriculture, much of it on family farms. Where bans have been imposed, it is not clear that they have been effective. Moehling (1999) shows that there is little evidence that child labor laws contributed to the dramatic decline in child labor in the nineteenth-century United States. This decline was driven instead by changes in technology, immigration, and the rise in the real wage.

Even if governments could effectively ban child labor, the consequences could be dire for those poor households (and their children) who are resorting to child labor out of desperation. These children are working to help the household make ends meet. An effective ban on child labor would make these households and these children worse off. Therefore any legal restrictions on child labor in developing countries should be focused on those cases

Cl

in
ta
th
cogu
of
ar
tec
we
the
age
in
tio
effe
the
the
wo
offgo
is
in
chil
grea
erin
makT
chal
on p
inter
diple
age
veryW
of ch
finan
from
well-
feasit
to th
mont
mark
nitud
Hc
Th

in which there is evidence of pathology, of parents or guardians who do not take into account the interest of the child. This is most likely in the case of the most odious of forms of child labor, including working under hazardous conditions or as bonded laborers.

A closely related issue is developed-country trade policy. Many have argued for an international labor standards policy that requires the elimination of child labor for access to developed country markets.⁶ In some cases, this argument is simply a smoke screen providing cover for standard-issue protectionism. However, it is often motivated by a genuine concern for the welfare of children in developing countries. If this is indeed the motivation, the implementation of trade sanctions to enforce an international standard against the use of child labor is likely to have perverse consequences. Except in unusual cases, which are discussed in Basu (1999, sec. 8), effective sanctions would make the families of child workers worse off. If sanctions are effective, they will generally have the consequence of lowering the price of the good produced with child labor. This reduction in the price would lower the value of workers to the employer, and thus lead to lower wages for child workers. Those children who remained in the labor force would be worse off because they would be paid less.

As a consequence of the lower wage, some children *will* stop working and go to school. That would seem to be a good thing! However, if child labor is indeed a means of coping with desperate poverty, families are sending children to work only when the current value of the income they earn is greater than the (discounted value) of the future benefits of education. Lowering the wage of child labor to induce the family to send the child to school makes the family worse off.

Trade sanctions are a particularly inappropriate tool for dealing with the challenge of child labor. As noted above, there is a legitimate case for bans on particular forms of child labor. If there is to be any developed-country intervention in these cases, it would be more appropriate in the form of diplomatic pressure and more general rewards and punishments to encourage government action to avoid the adverse effects of trade sanctions on the very children they are meant to help.

We have seen that dysfunctional financial markets are an important cause of child labor. Child labor would be dramatically reduced if parents could finance their children's exit from the labor force and entry into schooling from the increased future earnings of the child. Unfortunately, extremely well-functioning credit markets are required to make this kind of transaction feasible. The lag between the investment in child education and the return to that investment in the adult labor market is measured in decades, not months. There is little immediate prospect for improvements in financial markets accessible to the poor in developing countries of the order of magnitude required for such long-term transactions.

How, then, to reduce child labor?

The most effective way to draw children out of damaging work is to

encourage school attendance. One way of doing so would be to improve school quality, and therefore increase the gain from attending school. Handa (2002), for example, argues that school enrollment in Mozambique is quite sensitive to the number of trained teachers. This is an important tool that is available to reduce child labor. However, it has the significant disadvantage of influencing outcomes in the distant future, when the higher quality of schooling leads to higher wages as an adult. The influence of these changes on future outcomes of current decisions regarding work and schooling is scaled down by credit constraints and agency problems.

The most promising tool yet developed for reducing child labor is a targeted subsidy to families sending their children to school. In such a program, a grant is provided to the family of any child who is enrolled in school. The particular value of this intervention is that it addresses the root causes of child labor. It overcomes the problems associated with imperfect or nonexistent financial markets by balancing the current cost of moving a child out of the labor force and into school with a current grant. It addresses the main agency problem by providing current resources, thus reducing the importance of intergenerational transfers. For a priori reasons, then, we can expect subsidies for school enrollment to be a useful tool in the effort to reduce child labor. A number of such programs have been implemented in recent years.

The flagship program of this type is the innovative Progresa poverty program in Mexico (the name of the program has recently been changed to Oportunidades). Progresa provides mothers of enrolled students in rural Mexico with grants that have a value slightly less than the wage that would be earned by the child if she were working full-time. With remarkable foresight, the Progresa program was introduced (in 1998) in a randomized sequence. This randomization, combined with systematic data collection, makes it possible to measure with great confidence the impact of the program on both school enrollment and child labor force participation. Schultz (2004) estimates that the program has resulted in an increase in schooling of about two thirds of a year (from a baseline attainment of 6.8 years), and that child labor correspondingly falls. The most dramatic effects are for secondary school girls, whose broad labor force participation is estimated to drop by almost fifty percentage points upon enrollment in school.⁷

Inspired by the Progresa example, Nicaragua instituted the Red de Protección Social (RPS) in 2000. This program also provides grants for children aged seven to thirteen who attend school that are approximately as large (as a proportion of household expenditure) as those of Progresa. Maluccio and Flores (2004) analyzes the pilot phase of this program, in which (like Progresa) it was implemented in a randomly selected group of communities. Where it has been implemented, RPS has had a massive impact. Maluccio estimates that the effect of the program is to increase enrollment rates by twenty-two percentage points (from a base enrollment rate of 69%). The effects on child labor are large for older children (ten to thirteen years old);

Child La

younge
gram's
nine pThe
to the
averag
Wodo
rollm
pation

Ch

This
cumc
scho-
the c
mor
the
tries
isin

ACK

Rol
me
no
So
(gr
wy

N

(i
ii

younger children were unlikely to work even before the program. The program's impact was to reduce the proportion of children working by almost nine percentage points (from a base of 27%).

The Food for Education program in rural Bangladesh is similar in spirit to the other two programs. The monthly payment is smaller; 15–25% of average monthly earnings for working children. Nevertheless, Ravallion and Wodon (1999) estimate that the FFE program moved primary school enrollment from approximately 75% to over 90%. Child labor force participation dropped as well (by about 30% for boys and by about 20% for girls).

Child labor can effectively be reduced by subsidies for school enrollment. This tool dominates alternatives because it directly addresses the tragic circumstances that impel families to send their children to work instead of school. An effective subsidy program is not unreasonably expensive because the costs are tied to the low wages earned by child workers. Therefore, while more careful cost–benefit analyses should be completed on an urgent basis, the expansion of targeted education subsidies into areas of developing countries with high rates of child labor force participation is an extremely promising strategy.

ACKNOWLEDGMENTS I thank Amalavoyal Chari for valuable research assistance and Roland Benabou, Markus Goldstein, Barbara O'Brien, and Tavneet Suri for comments. I am grateful to the MacArthur Foundation Network on Inequality and Economic Performance, to the Pew Charitable Trust through its Project on Moral and Social Dimensions of Microeconomic Behavior in Communities, and to the NSF (grant 0079115) for financial support. Current versions of this essay can be found at www.econ.yale.edu/~udry/.

NOTES

1. See www.ilo.org/public/english/standards/norm/whatare/fundam/childpri.htm.
2. Ranjan (1999, 2001), Hazan and Berdugo (2002), and Baland and Robinson (2000) provide superb and simple theoretical models of the relationships between imperfect financial markets and child labor.
3. Manski (1995) is a superb treatment of the general problem.
4. When the wage does change, the “substitution effect” may dominate the “income effect” and child labor may decrease. Boozer and Suri (2001) show that agricultural shocks which decrease the productivity of labor are associated with declines in child labor and increases in school attendance in Ghana.
5. This is the collective household model, as in Chiappori (1988). Even this minimal assumption is somewhat controversial. I found evidence against it in Burkina Faso (Udry 1996) and, with Duflo, in Côte d'Ivoire (Duflo and Udry 2003).
6. Brown (2000) provides a very useful review of the large literature on this topic.
7. Preliminary cost–benefit analysis of the Progressa program by Schultz (2004) indicates a real rate of return of approximately 8%. The costs of the program are relatively easy to measure; Schultz limits the benefit calculation to the private market return to education. If there are important externalities or nonmarket returns to schooling, this is an underestimate.

BIBLIOGRAPHY

- Baland, Jean Marie, and James Robinson. "Is Child Labor Inefficient?" *Journal of Political Economy* 108 (4) (2000): 663–679.
- Basu, Kaushik. "Child Labor: Cause, Consequence and Cure, with Remarks on International Labor Standards." *Journal of Economic Literature* 37 (1999): 1083–1119.
- Becker, Gary. "A Theory of Social Interactions." *Journal of Political Economy* 82 (1974): 1063–1094.
- Beegle, Kathleen, Rajeev Dehejia, and Roberta Gatti. "Do Households Resort to Child Labor to Cope with Income Shocks?" Working paper, World Bank, 2002.
- Boozer, Michael, and Tavneet Suri. "Child Labor and Schooling Decisions in Ghana." Working paper, Yale University, 2001.
- Brown, Drusilla. "International Trade and Core Labor Standards: A Survey of the Recent Literature." Working paper, Tufts University, 2000.
- Chiappori, Pierre-Andre. "Rational Household Labor Supply." *Econometrica* 56 (1) (1988): 63–90.
- Duflo, Esther. "Grandmothers and Granddaughters: Old Age Pension and Intra-Household Allocation in South Africa." *World Bank Economic Review* 17 (1) (2003): 1–25.
- Duflo, Esther, and Christopher Udry. "Intrahousehold Resource Allocation in Côte d'Ivoire: Social Norms, Separate Accounts and Consumption Choices." Working paper, Yale University, 2003.
- Duryea, Suzanne, David Lam, and Deborah Levison. "Effects of Economic Shocks on Children's Employment and Schooling in Brazil." Population Studies Center Research Report 03–541, University of Michigan, Ann Arbor, 2003.
- Edmonds, Eric. "Is Child Labor Inefficient? Evidence from Large Cash Transfers." Working paper, Dartmouth College, Hanover, N.H., 2002.
- Edmonds, Eric, and Nina Pavcnik. "Does Globalization Increase Child Labor? Evidence from Vietnam." Working paper, Dartmouth College, Hanover, N.H., 2002.
- Fallon, Peter, and Zafiris Tzannatos. "Child Labor: Issues and Directions for the World Bank." Mimeo, World Bank, 1998.
- Grootaert, Christiaan, and Harry Patrinos. *The Policy Analysis of Child Labor: A Comparative Study*. New York: St. Martin's Press, 1999.
- Haddad, Laurence, John Hoddinott, and Harold Alderman, eds. *Intrahousehold Resource Allocation in Developing Countries: Models, Methods and Policy*. Baltimore: Johns Hopkins University Press, 1997.
- Handa, Sudhanshu. "Raising Primary School Enrollment in Developing Countries: The Relative Importance of Supply and Demand." *Journal of Development Economics* 66 (1) (2002): 103–128.
- Hazan, Moshe, and Binyamin Berdugo. "Child Labor, Fertility, and Economic Growth." *The Economic Journal* 112 (2002): 810–828.
- International Labour Office. *Every Child Counts: New Global Estimates on Child Labour*. Geneva: ILO, 2002.
- Jacoby, Hanan, and Emmanuel Skoufias. "Risk, Financial Markets and Human Capital in a Developing Country." *Review of Economic Studies* 64 (3) (1997): 311–335.
- Kruger, Diana. "The Effects of Coffee Production on Child Labor and Schooling in Brazil." Working Paper, University of Maryland, College Park, 2002.
- Moehling, Carolyn. "State Child Labor Laws and the Decline of Child Labor." *Explorations in Economic History* 36 (1999): 72–106.

- Maluccio, John, and Rafael Flores. "Impact Evaluation of a Conditional Cash Transfer Program: The Nicaraguan *Red de Proteccion Social*" Food Consumption and Nutrition Division Discussion Paper 184, International Food Policy Research Institute, Washington, D.C., 2004.
- Manski, Charles. *Identification Problems in the Social Sciences*. Cambridge: Harvard University Press, 1995.
- Pakistan Federal Bureau of Statistics. *Summary Results of Child Labour Survey in Pakistan*. Islamabad: Federal Bureau of Statistics, 1996.
- Psacharopoulos, George. "Child Labor Versus Educational Attainment: Some Evidence from Latin America." *Journal of Population Economics* 10 (1997): 377–386.
- Ranjan, Priya. "An Economic Analysis of Child Labor." *Economics Letters* 64 (1999): 99–105.
- Ranjan, Priya. "Credit Constraints and the Phenomenon of Child Labor." *Journal of Development Economics* 64 (2001): 81–102.
- Ravallion, Martin, and Quentin Wodon. "Does Child Labor Displace Schooling? Evidence on Behavioral Responses to an Enrollment Subsidy." Working paper, World Bank, 1999.
- Ray, Ranjan. "Child Labor, Child Schooling, and Their Interaction with Adult Labor: Empirical Evidence for Peru and Pakistan." *World Bank Economic Review* 14 (2) (2000): 347–367.
- Schultz, Paul. "Education Investment and Returns." In *Handbook of Development Economics*, Vol. 1. Edited by H. Chenery and T. N. Srinivasan. Amsterdam: North-Holland, 1995.
- Schultz, Paul. "School Subsidies for the Poor: Evaluating the Mexican Progreso Poverty Program." *Journal of Development Economics* 74 (1) (2004): 199–250.
- Udry, Christopher. "Gender, Agricultural Production and the Theory of the Household." *Journal of Political Economy* 104 (5) (1996): 1010–1047.