Transfers to Households with Children and Child Development*

Daniela Del Boca
University of Torino and CCA

Christopher Flinn
New York University and CCA

Matthew Wiswall
Arizona State University

December 2012
JEL Classification J13, D1
Keywords: Time Allocation; Child Development, conditional and unconditional cash transfer.

Abstract

In this paper we utilize a model of household investments in the cognitive development of children to explore the impact of various transfer policies on the distribution of child cognitive outcomes in target populations. We develop a cost criterion that can be used to compare the cost effectiveness of unrestricted, restricted, and conditional cash transfer systems, and find that conditional cash transfers are the most cost efficient way to attain any given gain in average child quality in a target population. Of course, this is only true if one uses efficiently designed cash transfer systems, and we are able to explore their design using our modeling framework.

1 Introduction

Over the past few decades there has been great interest in exploring the efficacy of various types of cash transfer programs aimed at supporting household investments in child development. A large body of research makes a case for there being large returns from investments in early childhood development. Empirical evidence of the strong impact of early investments by parents and others has motivated economic theories of skill formation in which early investments increase the productivity of later investments (e.g., Cunha et al., 2006). Children who are deprived of critical investments at young ages, because their

*Part of this research has been supported by a grant from the Institute of Human Development and Social Change at New York University and by a grant from the Collegio Carlo Alberto "Parental and Public Investments and Child Outcomes."

parents lack either resources or knowledge about parenting practices, are thought to be at high risk of negative consequences in later life (Heckman et al., 2006).

Although there is fairly broad consensus about the importance of investments in early childhood development, little is known about which specific types of interventions are likely to be most effective. Much of the research in the United States and other developed countries has focused on preschool interventions that provide young children with better environments outside of the home. The Perry Preschool Project and the Abecedarian Project have been particularly influential because they use a random assignment design and continue to follow the children well into their adult years. These studies demonstrate substantial positive effects of early environmental enrichment on a range of cognitive skills and behavioral traits, criminal behavior, school achievement, and job performance.[1]

In contrast, cash transfer programs that are designed to raise the living standards of poor families are becoming increasingly widespread across the developing world. While several studies have shown a positive correlation between unconditional cash transfers and child health, very little is known about the impact on child cognitive outcomes. Among the few studies on this topic, Paxton and Shady (2007) have analyzed the impact of a program (the Bono de Desarollo Humano, or BDH) which transfers cash to households in Ecuador conditional on child outcomes.[2] They find that the impact of these transfers (which are only $15 per month per family, but represent about a ten percent increase in family income for the average eligible family) is small but positive. Other studies (Fernald and Hidrobo, 2011) show that in rural areas, the program led to significantly better performance on the number of words in the child's vocabulary, and on the probability that the child was combining two or more words, while there were no significant effects on language development for children in urban areas.

There is a great deal of diversity across countries in how cash transfer programs are structured. In some cases, such as the Ecuadorean BDH or the South African Child Support Grant, cash transfers are means-tested but are not conditioned on specific parental behaviors, while in others, such as the Mexican Oportunidades program, receipt of money is conditional on specific behaviors, including enrolling children in school.

Conditional cash transfer programs (CCTs) have been increasingly implemented in developing countries since the 1990s. The first large-scale conditional cash transfer program was PROGRESA, which was launched in Mexico in 1997. More recently, CCT programs have been implemented extensively in other developing countries, such as Colombia, Nicaragua, Honduras, Brazil, Argentina, Ecuador, and Turkey, with the main goal being inducing parents to send their children to school instead of to work.

These programs provide low-income households with incentives to send their children to school by tying a cash transfer to school attendance and performance (Martinelli and

[1] It must be acknowledged that the extremely small sample sizes make the precise identification of the size of the effects of the Perry School experiment difficult to determine. See Heckman et al. (2011).

[2] The child outcome is measured using scores from the Spanish version of the Peabody Picture Vocabulary test, a widely used test on language ability.

Parker, 2003). The large empirical literature dedicated to the evaluation of CCTs shows that these programs boost school enrollment and decrease dropout rates (Skoufias and Parker (2001), Cardoso and Souza (2004), Attanasio et al. (2005), Behrman et al. (2005), Bourgignon et al (2003), Schady and Araujo (2006), Dubois et al. (2008), and Todd and Wolpin (2006)). However, the effects of CCTs on learning and school achievement are less widely investigated. Fernald et al. (2008) investigate the impact of conditional cash transfers on the cognitive and behavioral outcomes of children. They use the Mexican Oportunidades study, and exploit exogenous variation in the size of the transfers received by beneficiaries to conclude that larger transfers resulted in better cognitive development, possibly due to improvements in the quantity and quality of food consumed.

More recently, CCT programs have been implemented in developed countries such as the U.S. The first city which implemented and evaluated CCT programs was New York City (Opportunity NYC: Family Rewards). Other jurisdictions in the U.S. (the cities of Chicago, Illinois and Savannah, Georgia, and the state of California) and in the U.K. have actively considered, and in some cases piloted in small scale, CCTs as well. Rather than creating performance conditions based exclusively on "effort" (e.g., attending school), in New York City conditions were added that were based on the child's performance, especially educational performance on academic achievement tests (Aber and Rawlings, 2011). The cash incentives were conditional on activities and outcomes in children's education, as well as preventive health care, and parents' employment. After two years of the program, the study found positive effects on families' economic well-being and mixed effects on children's education, health care, and parents' employment. While the program did not affect significantly school outcomes for younger children, it substantially improved the achievement of older children (Morris et al 2012).

A common problem faced by these outcome-based CCTs is the design of the incentive system. The researchers implementing such a system must choose the set of agents to potentially receive rewards, performance targets, and reward sizes. The process of cognitive development and the nature of household interactions are most often not well enough understood to enable the policy maker to make informed choices regarding the design of the CCT, so that effective policies can only be learned through an extremely expensive process of trial and error.

The paper of the most direct methodological relevance to ours is Behrman et al. (2011). Here the authors report the results of their social experiment in Mexico, which is the Aligning Learning Incentives (ALI) program. The researchers allocated 88 Mexican high schools to three broad treatment groups and a control group. The treatments vary in terms of the agents targeted to receive the reward, and involve combinations of students, teachers, and administrators. Rewards are based on academic performance criteria, and the authors find that the treatment that rewards students, teachers, and school administrators in combination has the best performance outcome. The results here are striking and the evaluation analysis is extremely well done, but the results found are obviously conditional on the specific set of performance targets and reward levels utilized. It would be advantageous to combine the results from comprehensive field experiments, such as this one, with a behavioral model so as to learn what a more effective CCT design might look like.[3]

This paper builds on our earlier work, Del Boca, Flinn, and Wiswall (2011), to consider the design of "optimal" transfer systems in the context of our simple model of investments in child quality with heterogeneous household preferences, resource constraints, and a growth process for the child's cognitive development. While our model is extremely stylized, it does include all of these elements and hence we are able to consider the design of an effective, or even optimal, transfer system in the presence of extensive population heterogeneity in preferences and resources. We will not be going so far as to specify a complete social planning problem, but will confine our attention to the evaluation of the cost effectiveness of a restricted set of transfer mechanisms.

We consider the impact of three types of transfer programs on child development within the context of our model. The first are "unrestricted" transfers, in the sense that eligible households (with a child of a certain age) are transferred some fixed amount of dollars with no restriction on how it must be spent. The second are "restricted" transfers; in this case eligible households have to spend at least the amount that they receive in transfers on child investment goods. The final type of transfers considered are "conditional" ones. In this case, cash transfers are made to the household when it satisfies a performance target. While we explicate the setup for more general performance criteria, due to space constraints we look only at a cognitive growth criterion. The conditional transfer then is specified in terms of (i) the performance target and (ii) the level of the reward for meeting the target. We examine the tradeoffs in setting these two parameters and then compare the cost effectiveness of each of the three transfer systems.

The plan of the paper is as follows. In Section 2 we lay out the model structure from Del Boca et al. (2011). Section 3 provides a formal discussion of the very stylized mechanism design problem the social planner faces under the three types of transfer systems we consider. Section 4 contains the results from our simulations and the effectiveness comparisons across the three transfer types. Section 5 concludes.

2 Model Structure (DFW)

In this section we provide an overview of the model developed and estimated in DFW, which is important if the reader is to understand some of the mechanisms that underlie the results that are obtained below. Our discussion of the model is quite brief and succinct; the interested reader should consult the original paper for further details, as well as for model estimates and some comparative statics results.

[3]The model we exposit below does not include formal schooling, so school-based incentive systems are not considered below. In our current research, we are adapting the child cognitive ability production process to include formal schooling as an input. In this case, teachers and administrators will be able to alter educational quality when a CCT system is available to them.

The model is based on a set of assumptions that allow us to derive closed-form solutions to the household's dynamic optimization problem; it is the simple form of the life-cycle demand functions that allows us to include a relatively large number of endogenous variables in the model. In addition to assuming particular functional forms for the household's objective function and the child quality production technology, we assume that the household is not able to save or borrow. While we have estimated the model for the case of two-child families, for simplicity we only consider the one child case in our discussion here and in the analysis reported below.

The model in DFW is unique in the sense that it considers a large number of investment decisions, including the time investments of both the mother and father (the model only considers intact families), their labor supply decisions, and household consumption decisions. For the purposes of analyzing transfer policies, we believe that it is perhaps the best (estimated) model of the child development process currently available. The model of Bernal (2008) examines maternal choices of labor supply and child care for young children in a model of cognitive ability development, and while allowing for endogenous wage growth (not considered in our model), it neglects the role of the father in providing child investments and considers all of the mother's time away from work as child investment time. We find that mothers consume substantial amounts of leisure away from work, and that the father's time in child investment is almost as valuable as the mother's at certain stages of the development process. We think that it is important to account for these decisions in a model of the child development.

The work reported in a sequence of papers by Heckman and coauthors (e.g., Cunha and Heckman (2008), Cunha et al. (2010)) has been extremely influential in establishing the importance of noncognitive skills in adolescent and adult social and economic outcomes and the critical role of early intervention in mediating negative environmental influences in the lives of many disadvantaged children. The framework used in these empirical analyses are in several ways more general than is the data generating process implied by DFW. The limitation of that work for conducting the sorts of policy experiments considered in this paper is that the decision rules of the household are not explicitly modeled. It is for these reasons that the DFW setup, though extremely stylized, is perhaps the best one available to analyze counterfactual childhood cognitive development policies.

2.1 Timing and Preferences

The model begins with the birth of a child. The household makes decisions in each period of a child's life (or, more accurately, over the development period that we model), where the child's age is indexed by t. Parents make investments in child quality from the first period of the child's life, $t = 1$, through the last developmental period, M. At this "terminal" point (from the perspective of the parents' investment in the child), the child has reached adulthood and adult outcomes depending (in part) on the level of child quality obtained at this point.[4]

In each period, the household makes seven choices: hours of work for each parent: h_{1t} (mother) and h_{2t} (father); time spent in "active" child care for each parent: $\tau_{1t}(a)$ (mother) and $\tau_{2t}(a)$ (father); time spent in "passive" child care by each parent: $\tau_{1t}(p)$ and $\tau_{2t}(p)$; and expenditures on "child" goods, e_t. Household utility in period t is a function of each parent's hours of leisure, l_{1t} for the mother and l_{2t} for the father, the level of a consumption good produced by the household, c_t, and the level of their child's quality, k_t. We assume a Cobb-Douglas household utility function and restrict the preference parameters to be stable over time:

$$u(l_{1t}, l_{2t}, c_t, k_t) = \alpha_1 \ln l_{1t} + \alpha_2 \ln l_{2t} + \alpha_3 \ln c_t + \alpha_4 \ln k_t, \tag{1}$$

where $\sum_j \alpha_j = 1$. In the empirical implementation of the model, we allow heterogeneity in the parameter vector α across households.

Before we proceed to the description of the production technology, note that time with children is purely an investment in child quality. There is no direct utility from time with children, i.e. "enjoyment" of time with children or some effort cost of this time. A model with these elements would be one where time investments had multiple outputs (both utility and child quality). In our model, the value of the child to the household is captured through the enjoyment of child quality, which depends on all investments, from both parents and non-time expenditures.

2.2 Child Quality Production

Age $t+1$ child quality is produced by the current level of child quality, k_t, parental time investments in the child of the active and passive kind, and expenditures on the child, all of which are made when the child is age t. We assume a Cobb-Douglas form for the child quality technology:

$$\begin{aligned} k_{t+1} &= f_t(k_t, \tau_{1t}(a), \tau_{2t}(a), \tau_{1t}(p), \tau_{2t}(p), e_t) \\ &= R_t \tau_{1,t}(a)^{\delta_{1,t}(a)} \tau_{2,t}(a)^{\delta_{2,t}(a)} \tau_{1,t}(p)^{\delta_{1,t}(p)} \tau_{2,t}(p)^{\delta_{2,t}(p)} e_t^{\delta_{3,t}} k_t^{\delta_{4,t}}, \end{aligned} \tag{2}$$

where $R_t > 0$ is the scaling factor known as total factor productivity, or TFP.

While the Cobb-Douglas form restricts the substitution possibilities, we allow the productivities of the various inputs to vary over the age of the child. This allows us to capture the important insights in the economics and child development literature that the marginal

[4] The terminal date M needs not correspond to the end of the investment period in the child. In a more elaborate model of child development, it may correspond to the end of a particular developmental stage, with the final value of child quality in the current stage of development serving as an initial condition into the next stage of development, which may be characterized by very different production technologies. While we have not pursued such an approach in this paper, it is a subject of our on-going research.

productivity of inputs varies over the stages of child development (for a useful survey, see Heckman and Masterov (2007)). As written in (2), the production technology is deterministic assuming knowledge of the $\{R_t\}_{t=1}^M$ and $\{\delta_t\}_{t=1}^M$ sequences.

2.3 Dynamic Problem

Given wage offers and the current level of child quality, parents optimally choose their labor supply and child inputs to maximize expected lifetime discounted utility. The value function for the household at period t is then

$$\begin{aligned} V_t(S_t) = & \max_{l_{1t}, l_{2t}, \tau_{1t}(a), \tau_{2t}(a), \tau_{1t}(p), \tau_{2t}(p), e_t} u(l_{1t}, l_{2t}, c_t, k_t) + \beta E_t V_{t+1}(S_{t+1}), \\ & \text{s.t. } T = l_{jt} + h_{jt} + \tau_{jt}(a) + \tau_{jt}(p),\ j = 1, 2 \\ & c_t + e_t = w_{1t} h_{1t} + w_{2t} h_{2t} + I_t \end{aligned}$$

where the vector of state variables S_t consist of the current level of child quality, the wage offers to the parents, and nonlabor income,

$$S_t = (k_t\ w_{1t}\ w_{2t}\ I_t), \tag{3}$$

β ($\in [0, 1)$) is the discount factor, and E_t denotes the conditional expectation operator with respect to the period t information set. The conditional expectation is taken with respect to the random variables appearing in the household's period $t+1$ problem, which include wages for both parents, household nonlabor income, and possibly R_{t+1}. The state variable vector at the birth of the child are the initial conditions of the problem, $S_1 = (k_1\ w_{11}\ w_{21}\ I_1)$.

The constraint set faced by the household in period t consists of time and market good expenditures restrictions. We assume that each parent has a time endowment of T hours, and that this time is allocated between leisure, market labor supply, active time spent with the child, and passive time spent with the child. The last constraint is the expenditure constraint, and its form follows from our assumption that there is no saving and borrowing and that the prices of c_t and e_t are 1 in every period.

2.4 Terminal Value

We think about the child development process as lasting for M periods, and resulting in a "final" child quality level of k_{M+1}. Parental investments in child quality are limited to the first M period's of the child's life during the development period we study. We think of the child quality level k_{M+1} as an initial condition into a second stage of the child development process, one that may (and most surely does) include investment by the child in their own cognitive development, savings by parents and the child (possibly) for college costs, etc. Since the only truly dynamic process in our model is that of the child's cognitive development, the only "carry over" from the development stage we model is the child quality level at the beginning of the new development stage, k_{M+1}. We assume that the value of k_{M+1} at the beginning of the next development stage (i.e., period $M+1$) is given by $\psi\alpha_4 \ln k_{M+1}$, where ψ is a parameter to be estimated. Therefore we can write the period M optimization problem as

$$
\begin{aligned}
&V_M(w_{1M}, w_{2M}, I_M, k_M) = \\
&\max_{l_{1,M}, l_{2,M}, \tau_{1,M}(a), \tau_{1,M}(p), \tau_{2,M}(a), \tau_{2,M}(p), e_M} \alpha_1 \ln l_{1M} + \alpha_2 \ln l_{2M} + \alpha_3 \ln c_M + \alpha_4 \ln k_M \\
&+\beta\psi\{\delta_{1,M}(a) \ln \tau_{1,M}(a) + \delta_{2,M}(a) \ln \tau_{2,M}(a) + \delta_{1,M}(p) \ln \tau_{1,M}(p) \\
&+\delta_{2,M}(p) \ln \tau_{2.M}(p) + \delta_{3,M} \ln e_M + \delta_{4,M} \ln k_M\}
\end{aligned}
$$

2.5 Model Solution

As is clear from the nature of the production technology, there are never any corner solutions to the household input choice problem during the investment period.[5] However, we do allow for corner solutions in labor supply as labor supply for either or both parents may be 0 in a given period. We can write the conditional factor demands for child inputs, where we are conditioning on labor supply choices and nonlabor income, as

$$
\begin{aligned}
\tau^*_{1,t}(a) &= (T - h_{1t}) \frac{\varphi_{1,t}(a)}{\alpha_1 + \varphi_{1,t}(a) + \varphi_{1,t}(p)} && (4) \\
\tau^*_{2,t}(a) &= (T - h_{2t}) \frac{\varphi_{2,t}(a)}{\alpha_2 + \varphi_{2,t}(a) + \varphi_{2,t}(p)} && (5) \\
\tau^*_{1,t}(p) &= (T - h_{1t}) \frac{\varphi_{1,t}(p)}{\alpha_1 + \varphi_{1,t}(a) + \varphi_{1,t}(p)} && (6) \\
\tau^*_{2,t}(p) &= (T - h_{2t}) \frac{\varphi_{2,t}(p)}{\alpha_2 + \varphi_{2,t}(a) + \varphi_{2,t}(p)} && (7) \\
e^*_t &= (w_{1t} h_{1t} + w_{2t} h_{2t} + I_t) \frac{\varphi_{3,t}}{\alpha_3 + \varphi_{3,t}} && (8)
\end{aligned}
$$

where

$$
\begin{aligned}
\varphi_{l,t}(\xi) &= \beta\delta_{l,t}(\xi)\eta_{t+1}, \; l = 1,2; \; \xi = a, p, \\
\varphi_{3,t} &= \beta\delta_{3,t}\eta_{t+1}.
\end{aligned}
$$

The sequence $\{\eta_t\}_{t=1}^{M+1}$ is defined (backwards-) recursively as

[5] If any factor is set at 0, then child quality will be 0 in all subsequent periods, and household utility diverges to $-\infty$ as $k \to 0$ whenever $\alpha_4 > 0$.

$$
\begin{aligned}
\eta_{M+1} &= \psi\alpha_4 \\
\eta_M &= \alpha_4 + \beta\delta_{4,M}\eta_{M+1} \\
&\vdots \\
\eta_t &= \alpha_4 + \beta\delta_{4,t}\eta_{t+1} \\
&\vdots \\
\eta_1 &= \alpha_4 + \beta\delta_{4,2}\eta_2.
\end{aligned}
$$

where η_t represents the period t marginal utility of (log) child quality to the household: $\eta_t = \partial V_t(S_t)/\partial \ln k_t$. η_t reflects both the present period flow marginal utility of (log) child quality to the household, given by α_4, and the discounted value of child quality to future utility. The latter value of current child quality depends on the discount rate and the technologically determined productivity of the current stock of child quality in producing future child quality, given by the time varying parameter $\delta_{4,t}$.

The solution to the spousal labor supplies problem in period t also has a simple form. Define two "latent" labor supply variables in period t by

$$
\begin{aligned}
\hat{h}_{1t} &= \frac{A_{1t} - A_{2t}B_{1t}}{1 - A_{2t}B_{2t}} \\
\hat{h}_{2t} &= \frac{B_{1t} - B_{2t}A_{1t}}{1 - A_{2t}B_{2t}},
\end{aligned}
\tag{9}
$$

where

$$
\begin{aligned}
A_{1t} &= \frac{w_{1t}T(\alpha_3 + \varphi_{3,t}) - (\alpha_1 + \varphi_{1,t}(a) + \varphi_{1,t}(p))I_t}{w_{1t}(\alpha_1 + \alpha_3 + \varphi_{1,t}(a) + \varphi_{1,t}(p) + \varphi_{3,t})} \\
A_{2t} &= \frac{w_{2t}(\alpha_1 + \varphi_{1,t}(a) + \varphi_{1,t}(p))}{w_{1t}(\alpha_1 + \alpha_3 + \varphi_{1,t}(a) + \varphi_{1,t}(p) + \varphi_{3,t})} \\
B_{1t} &= \frac{w_{2t}T(\alpha_3 + \varphi_{3,t}) - (\alpha_2 + \varphi_{2,t}(a) + \varphi_{2,t}(p))I_t}{w_{2t}(\alpha_2 + \alpha_3 + \varphi_{2,t}(a) + \varphi_{2,t}(p) + \varphi_{3,t})} \\
B_{2t} &= \frac{w_{1t}(\alpha_2 + \varphi_{2,t}(a) + \varphi_{2,t}(p))}{w_{2t}(\alpha_2 + \alpha_3 + \varphi_{2,t}(a) + \varphi_{2,t}(p) + \varphi_{3,t})}.
\end{aligned}
$$

Given these latent labor supplies, we can define the actual optimal hour choices that satisfy the rationing constraint on the time allocations of the parents. If the latent labor supplies on the right hand sides are set to zero, it is apparent that the condition required for the conditional latent labor supplies to both be 0 is

$$
(h^*_{1t} = 0, h^*_{2t} = 0) \Leftrightarrow A_{1t} \leq 0 \text{ and } B_{1t} \leq 0.
$$

If both of these intercept terms are nonpositive, then the household supplies no time to the market. For this to be the case, it is necessary that the household's nonlabor income be strictly positive.

Going back to the "full" solutions to the model given in (9), if both of the solutions are positive, then both satisfy the time allocation constraints, and these are the solutions to the household optimization problem. If the latent labor supply of parent 1 is positive and that of parent 2 is negative, then ($h^*_{1t} = A_{1t}, h^*_{2t} = 0$), while if the situation is reversed, the solution is ($h^*_{1t} = 0, h^*_{2t} = B_{1t}$). In summary, optimal labor supplies are

$$(h^*_{1t}, h^*_{2t}) = \begin{cases} (0,0) & \text{if} \quad A_{1t} \leq 0 \text{ and } B_{1t} \leq 0 \\ (A_{1t}, 0) & \text{if} \quad A_{1t} - A_{2t}B_{1t} > 0 \text{ and } B_{1t} - B_{2t}A_{1t} < 0 \\ (0, B_{1t}) & \text{if} \quad A_{1t} - A_{2t}B_{1t} < 0 \text{ and } B_{1t} - B_{2t}A_{1t} > 0 \\ (\hat{h}_{1t}, \hat{h}_{2t}) & \text{if} \quad A_{1t} - A_{2t}B_{1t} \geq 0 \text{ and } B_{1t} - B_{2t}A_{1t} \geq 0 \end{cases}$$

Using these optimal labor supply choices, the investment decisions are determined using (4), (5), (6), (7), and (8) after substituting h^*_{1t} and h^*_{2t} into the functions.

2.6 Characteristics of Decision Rules

We conclude this section by discussing a few characteristics of the decision rules derived under our modeling assumptions. Most notable is the fact that our functional form assumptions result in decision rules that are independent of the current child quality state, though the decisions are a function of the parameters of the child quality production process. Child quality remains a state variable in the problem since it enters the utility function of the household in every period. The lack of dependence of investment and labor supply decisions on child quality levels greatly simplifies the computational burden of solving the model, enabling us to find closed-form solutions for all seven endogenous variables. Even though the functional form assumptions are restrictive, it is not necessary to assume temporal invariance of either the child quality production function or of household preferences. For purposes of estimation where identification of time varying preferences may be difficult, we have assumed time-invariant household preferences, but in principle there is no restriction on the manner in which the Cobb-Douglas utility function could vary over time. This makes the specification sufficiently flexible to fit most patterns in the data while preserving its very attractive computational properties.

The model easily accommodates exogenous variation in wages and nonlabor incomes over the period of the development process, as well as temporal variability in total factor productivity in the production process. The model solutions are invariant with respect to these sources of randomness, given our assumption of no borrowing and saving. Even though the model is stylized and quite parsimonious, we found that it was able to adequately capture the main features of the household labor supply, child investment, and child development processes both across households and over the development period.

2.7 Estimates

We briefly review some features of the estimated model from DFW and the insights it is capable of yielding in terms of the impact of changes in the household's resource constraints on its behavior and child outcomes.

Table 1 contains estimates of the first two moments of the distribution of preference parameters in the population of households with one child. Recall that the preference weights are normalized to sum to 1 for each household, and that parent 1 is the mother, parent 2 is the father. The first thing to note is that households, on averages, attach the largest utility weight to their child's cognitive ability. However, the weight attached is less than the sum of the weights attached to parental leisures, and is only about 37 percent greater than the weight attached to household consumption. Thus while household welfare is strongly linked to child quality, it is by no means the only determinant of household utility.

The last column of the table contains estimates of the coefficient of variation associated with the marginal distributions of the utility weights in this population. Here we see that there is substantial heterogeneity in preference parameters across households, particularly the value of the mother's leisure (0.619) and the value of the child's cognitive ability (0.568). These estimates indicate that population heterogeneity in tastes could be an important concern when designing policies to positively impact the distribution of child cognitive outcomes in the population.

Most interventions that are discussed in the literature, carried out in social experiments, or that are analyzed in this paper, involve monetary or in-kind transfers of child investment goods to the household. The impact of such transfers hinges on how important these types of investment goods are in the development of child cognitive ability. The evidence obtained in DFW indicates that the goods expenditures on children have productive values that are dominated by parental time investments, particularly in the early stages of the development process. As discussed above, we allow the productivity of the five forms of investment, parental time of the active and passive variety and goods investments, to vary over the developmental period. Figures 1 and 2 display the age-profile of the productivity parameters. Recall that the dynamic production technology has age $t+1$ cognitive ability being a function of these five inputs and child cognitive ability from period t. We see that the productivity of active time investments of mothers and fathers are largest at early ages, and that the father is an important actor in the child's development through the development period. When children enter into formal schooling, the active and passive time productivities of both parents are essentially identical and are small.

In Figure 2 we see that the productivity of investment goods are low when the child is born (0.05). Over time, these investment goods become much more productive, and at the end of the development process have a productivity of 0.18. In the same figure, we see that the child's cognitive outcomes become more difficult to change over time, as evidenced by the increasing importance of the previous period's ability in determining current ability.

These figures show that investment goods only become important determinants of cognitive outcomes late in the development period, and that it becomes more difficult to alter the child's cognitive ability as he or she ages, a point often made in recent research on the subject using other analytical frameworks.

We conclude this review of DFW with a brief consideration of Table 2, which traces how changes in the parents' wages impact household choices and cognitive ability outcomes. While child development policies do not typically target the wages of the parents, we present this table to indicate the complex patterns of behavioral responses that changes in the resource constraint of the household engender. These types of patterns will also exist in the analyses conducted below, in which the nonlabor income stream of the household is changed through the receipt of a transfer, or the potential receipt of one in the case of conditional (on household behavior) transfer programs.

For all simulation paths used to estimate the model, the wage paths of a given parent were increased in each period by a constant percentage, and elasticities were computed for each of the choices and outcomes listed. For the case of child ability outcomes, the elasticity is computed using the end of development period cognitive ability of the child. For other cases, such as the mother's labor supply, the response is measured in terms of average labor supply over the development period.

We see that wage increases have little effect on the age 16 cognitive ability of the child, and we can also see what this is the case. When the wage profile of the mother shifts up, she supplies more time to the market, and reduces her leisure and contact time with the child. The father increases his leisure and time with the child, and since the value of the father's time in producing child quality is not far below that of the mother's, except for very small children, there is no significant loss in the value of the time invested in the child through this substitution pattern. Of course, the wealth of the household has increased, and consequently more is spent on child investment goods, though as we have seen, these have little impact on child quality except at the end of the development period. A similar type of pattern emerges when we look at responses to increases in the father's wage profile.

Taken together, these findings suggest that transfers to households with children intended to increase child ability outcomes will be diverted to a number of other uses unless the policy-maker can restrict, in some manner, the household's use of these transfers, or can make the receipt of these transfers conditional on the household investing sufficient time and money to satisfy performance criteria set by the policy-maker. They also suggest that the impacts of transfers of whatever type will vary with the developmental age at which they are received and will be heterogeneous in the population.

3 Monetary Incentives and Child Development

In this section we consider the impact of three types of transfer programs on child development within the context of our model. For simplicity, we only consider the case of one-child households and in performing the numerical evaluations we use point estimates of the primitive parameters from that specification.[6]

The programs we consider vary in the degree to which the monetary transfers to households with children of a certain age t are designed to specifically increase the child's cognitive performance. They also differ in their costs of implementation, a factor not often considered when exploring the design of efficient social policies. Our goal will be to compare the costs of implementing these policies to achieve a given gain in the aggregate cognitive ability stock of the population of a subpopulation of children when transitioning from age t to age $t+1$. Our benchmark will be the cost to achieve a Z^* point gain in the average cognitive ability of a group of age t children relative to what it would have been in the absence of the program.

3.1 Targeting Sub-Populations

In the discussion that follows, it is useful to describe a household at time t in terms of the its characteristics ω_t, which include those that directly influence child development outcomes in period $t+1$ and those that do not. The state variable vector S_t that was defined in (3) is included in ω_t, but so are other characteristics of the household that do not directly influence investments in period t and outcomes in period $t+1$. We will denote all of these other characteristics by ν_t; some of these characteristics are time-varying (such as the neighborhood in which the household lives), while others are time-invariant (such as the racial composition of the household). Then $\omega_t = (S_t' \ \nu_t')'$.

We will speak of the targeting of transfers a number of times below, and by this we will mean that a subset of the population of interest to the social planner will be given access to the transfer program. The planner may use targeting of this form to increase the cost effectiveness of the transfer program. For example, while the planner might be interested in improving cognitive outcomes in the entire black population, to increase the cost effectiveness of the program he may make the transfer payments only available to black households in which the father and mother have wages under some amount $\hat{w}$. If the entire population of interest to the household belongs to the set Ω, so that $\omega \in \Omega$ is a household in this population, then we denote the subset targeted for receiving transfers as $\zeta \subseteq \Omega$. As is indicated by the notation, the targeted population can consist of the entire population of interest, and most often will, in the simulation results reported below.

We want to draw attention to the fact that the targeted group need not be defined *directly* in terms of the state variables of the model that determine investments and (in a

[6] The analysis reported here is supposed to be merely suggestive in nature. Since we are able to construct sampling distributions of the estimators using bootstrap methods, a serious evaluation of policy should include the evaluation of the effects of the policy under a large number of parameter vector draws from the sampling distribution of the estimator. Considerable computational effort is required to do so and so given the large number of trial programs we consider, we limit our attention to the results associated with the point estimates only.

stochastic manner) cognitive outcomes. In the model, the state variables that determine investment are the sequence of wage draws for the parents, nonlabor income realizations, and the preference parameters that characterize the household. The targeted group need not be defined in terms of these characteristics. Even when being concerned about improving cognitive ability outcomes in the entire population, the planner may limit the availability of transfers to a subpopulation not explicitly defined in terms of income draws or taste parameters. For example, the planner may want to offer transfers only to a set of individuals defined by values of nonlabor income and wage draws, but these are difficult to truthfully elicit. If the target is set in terms of total household income, families with certain wage draws may alter their labor supply decisions so as to qualify for transfers. Households may also have an incentive to hide nonlabor income to qualify for transfers. For all of these reasons, sometimes targeting based on predetermined demographic characteristics may be preferable due to the low cost of determining eligibility and the fact that households cannot strategically alter their choices so as to qualify for the program.

We can make this point more formally by using the definitions of ω given above. Let there be a targeted group for transfers defined in terms of ζ, where ζ is a subset of characteristics included in ν_t. Then a necessary condition for this targeting to be a viable means to affect household behavior it must be the case that

$$F_{S_t|\zeta}(\cdot|\zeta) \neq F_{S_t}(\cdot), \tag{10}$$

where $F_{S_t|\zeta}(\cdot|\zeta)$ is the conditional distribution of state variables for households characterized by ζ and $F_{S_t}(\cdot)$ is the marginal distribution of state variables. (10) expresses that the state variables and the criterion set ζ cannot be independently distributed. For example, defining the subpopulation targeted for transfers to be African-American households is potentially useful for improving the the population distribution of cognitive outcomes in period $t+1$ if African-American households have a different distribution of the state variables S_t than do non-African-American households. We know that this is the case, so targeting this subpopulation (of for that matter, non-African-American households), is potentially useful in terms of lowering costs associated with obtaining an improvement of Z^* in the entire population.

3.2 Criteria for Policy Evaluation

In terms of defining the improvement target Z^*, there are clearly an enormous number of potential choices. The target could be set in terms of: raising the cognitive ability outcome of quantile j in the distribution from its status quo value; decreasing the variance of the period $t+1$ cognitive outcome distribution; increasing the average cognitive ability in the target population; and many others. Once again, for simplicity and due to space limitations, we will define the planner's interests solely in terms of improving average quality in the target population of interest. We define the average cognitive ability of children of age $t+1$ in the population of interest as

$$\int k_{t+1}(s_t)dF_{S_t}(s_t)$$

in the absence of a transfer program.

We will generally describe a transfer program by $\varphi(\zeta)$. Note that the program occurs in period t and affects the period $t+1$ production of cognitive ability k_{t+1}. Then in the presence of the transfer program, average age $t+1$ cognitive abilities in the population of interest are

$$\int k_{t+1}(s_t, \varphi(\zeta))dF_{S_t}(s_t).$$

Then we define the targeted improvement on sub-population characterized by ζ as

$$\frac{\int k_{t+1}(s_t, \varphi(\zeta))dF_{S_t}(s_t)}{\int k_{t+1}(s_t)dF_{S_t}(s_t)}, \tag{11}$$

Note that throughout we have implicitly assumed that there are no spillover effects of the transfer policy from the "treated" sub-population to the "untreated" sub-population. The sub-population not receiving the transfer then remains at the baseline level of cognitive ability.

The canonical transfer program design problem indexed by the level of improvement Z^* can then be stated as

$$\begin{aligned} &\min_{\varphi(\zeta)} C(\varphi(\zeta)) \\ \text{subject to } Z^* \quad &= \quad \frac{\int k_{t+1}(s_t, \varphi(\zeta))dF_{S_t}(s_t)}{\int k_{t+1}(s_t)dF_{S_t}(s_t)}, \end{aligned}$$

where $C(\varphi(\zeta))$ are the total monetary costs (solely in terms of transfers made) associated with the transfer policy $\varphi(\zeta)$. We do not consider administrative costs associated with any of the programs we evaluate through simulation, but they are all program types that are used and which are not too "exotic."

Throughout our discussion and the policy simulations, we will only consider the case of short-run program interventions. In this situation, an unanticipated program is announced when parents have a child of age t, the parents adjust their behavior accordingly, and child outcomes are observed when the child is aged $t+1$. We do not consider the case in which any of the programs we describe are continued for more than one period. This is primarily due to our desire to keep the analysis simple and the set of policy options small. There are other reasons as well. When the reward system is defined over a sequence of outcomes of the child, this raises the possibility that the agents[7] receiving the award will act in a strategic

[7]In our model the only agents actively contributing to the child's cognitive development are the parents. More generally, these agents include teachers, relatives, other students, and school administrators.

manner. For example, when the reward is based on the sequence of improvements in child test scores over some period of the development cycle, the agents may have an incentive to manipulate the sequence of investments so as to produce the highest possible payoff.[8] Generally speaking, a payoff that is a function of a sequence of outcome realizations will have a complex effect on the dynamic pattern of investments. For these reasons, we ignore such programs in the current analysis, though our modeling framework can be used to analyze such cases after appropriate adjustments.

3.3 Unrestricted Transfers

In the unrestricted transfer case, we determine the amount of money that would have to be transferred to all households of children of age t in the subpopulation ζ to increase the average level of child cognitive ability by Z^* in comparison with what it would have been in the absence of the program. By "unrestricted" we mean that an eligible household in group ζ is transferred some amount of dollars with no restriction on how it must be spent. From our behavioral model, we know that such a transfer will be treated as nonlabor income by the household, and a portion of that additional income will be spent on child investment goods and increased parental contact time with the child, but some of the transfer will be "taxed away" by the household and spent on household consumption goods and increased leisure of the parents. The unrestricted lump sum transfer that we are considering here has as a benefit the simplicity of the transfer policy and the lack of enforcement required to implement it.

It need not be the case that the unrestricted money transfers are the same to all households in the target population for transfer receipt, ζ. However, for simplicity, we will only consider the case in which a common dollar transfer amount is made to all households in ζ. To distinguish this transfer amount from the next two we consider, we will denote it by φ_U. The impact of this transfer policy on the behavior of the household is quite transparent as we have just seen, since any household in the transfer target population ζ sees their nonlabor income increase at child age t from I_t to $I_t + \varphi_U$. The cost of this type of transfer is $C(\varphi_U(\zeta)) = \#(\zeta)\varphi_U$, where $\#(\zeta)$ is the size of the subpopulation ζ. Then the optimal unconditional transfer policy is defined by

$$\begin{aligned} & \min_{\varphi_U,\zeta} \#(\zeta)\varphi_U \\ \text{subject to } Z^* &= \frac{\int k_{t+1}(s_t, \varphi_U(\zeta)) dF_{S_t}(s_t)}{\int k_{t+1}(s_t) dF_{S_t}(s_t)}. \end{aligned}$$

[8] Macartney (2012) contains a nice theoretical and empirical illustration of ratchet effects in a pay for performance mechanism introduced in schools in North Carolina.

3.4 Restricted Transfers

A restricted transfer is the same as an unrestricted one except that the institutional agent requires the household to spend at least the amount φ_R on the child. In order to enforce such a condition, the social planner would have to verify the child good expenditures of the household. An equivalent transfer policy is to make transfers in-kind, that is, in the form of child investment goods. If the planner makes such a transfer, then the amount φ_R has no value outside of its use in the child production technology, and so the constraint is trivially satisfied.[9]

This kind of transfer acts as a constraint on the allocation decisions of some households. Say that prior to notification of this transfer, the household plans an expenditure of e_t. Then if $e_t \geq \varphi_R$, the receipt of an amount of goods φ_R that are perfect substitutes for already planned child goods expenditures means that the value of the transfer amount from the point of view of household utility maximization is the same as that of receiving an unrestricted transfer of φ_R directly. For households that anticipated spending far less than φ_R on the child, the transfer does distort the investment decisions of the household in favor of expenditures on child investment goods.[10] All households receiving such a transfer must attain a higher utility level, though the utility level can be no higher than it would be if the same amount were unrestricted.

If we abstract from the cost of making the restricted transfers, which could include monitoring household purchases or making deliveries of in-kind transfers, then the cost is the same as in the unrestricted case, $\#(\zeta)\varphi_R$. The planner's problem is

$$\min_{\varphi_R,\zeta} \#(\zeta)\varphi_R$$
$$\text{subject to } Z^* = \frac{\int k_{t+1}(s_t, \varphi_R(\zeta))dF_{S_t}(s_t)}{\int k_{t+1}(s_t)dF_{S_t}(s_t)}.$$

3.5 Conditional Transfers[11]

In the conditional transfer case (for symmetry we will refer to this case as conditional transfers as opposed to conditional cash transfers (CCT)), the household only receives the transfer amount if it qualifies on the basis of some performance criteria. In many CCT

[9]This assumes, of course, that there is no secondary market on which these child investment goods can be exchanged for money or some other good other than a child investment good.

[10]For households who would spend less than φ_R on the investment good before the transfer but would spend more than φ_R on the investment good if the transfer was unrestricted, the restriction on expenditures is not binding. For example, assume that the transfer amount is $300. If a household would spend $299 on the child good in the absence of a transfer and would spend $320 if the transfer of $300 was unrestricted, the receipt of the transfer does not distort investment decisions.

[11]For notational simplicity, in this section we do not explicitly consider the dependence of the policy on the subset of the population eligible to participate in the conditional transfer program, ζ. In this way we can focus on the other two elements of the conditional transfer policy, the size of the award threshold ξ and the payment φ_C.

programs, such as PROGRESA, the transfer is received if the household engages in certain behaviors, such as sending their child to school or taking the child to see a doctor. We will look, instead, for conditional transfers that are made when the child's cognitive ability in periods t and $t+1$ exceeds a preset threshold ξ. The generic form of the cash transfer φ_C to a household in the subpopulation ζ eligible to receive these transfers is

$$\tilde{I}_{t+1} = \begin{matrix} I_{t+1} + \varphi_C & \text{if} & \rho(k_t, k_{t+1}) \geq \xi \\ I_{t+1} & \text{if} & \rho(k, k_{t+1}) < \xi \end{matrix},$$

where ρ is a function that defines the performance criterion.

We consider two cases. In the first, we have

$$\rho_1(k_t, k_{t+1}) = k_{t+1}.$$

In this case, there is an absolute standard that all children must reach for the household to receive the transfer φ_C in period $t+1$. In the second case, the planner utilizes more of a "value-added" criterion, and sets

$$\rho_2(k_t, k_{t+1}) = k_{t+1}/k_t.$$

In the first case, then, the reward criterion is set solely in terms of levels, and in the second it is set in terms of growth. Even when ρ_2 criterion is used, however, we continue to assume that the planner's objective is to increase the average level of cognitive ability in the population of interest relative to the baseline level. That is, we do not assume that the planner's objective is formulated in terms of the growth of cognitive ability between t and $t+1$, even if the the award criterion is formulated in terms of the size of the improvement in the child's cognitive ability.

All programs such as this generate a certain amount of rents that are unevenly distributed in the population due to the fact that the program cannot condition on all relevant household characteristics, such as preferences, production technologies (which are assumed homogeneous in our case but that can vary with the age of the child), and household resources, that are included in ω. In the absence of the program, the optimal choices of a household of type ω will result in a path of child qualities $(k_1, k_2, ..., k_{M+1})(\omega)$. In the case of the criterion ρ_1, if $k_{t+1}(\omega) \geq \xi$, then the household satisfies the award criterion and obtains the transfer φ_C even without modifying its behavior by spending more household resources on child cognitive development in period t. The same holds for the criterion ρ_2, when under the baseline $k_{t+1}(\omega)/k_t(\omega) \geq \xi$.

Consider the simpler criterion, ρ_1. When a household eligible for such an award has $k_{t+1}(\omega) \geq \xi$, this is a pure rent in the sense that nonlabor income is increased in period $t+1$ without the household modifying its behavior. For any level of cognitive ability ξ, there will exist a set of households that will earn the reward with no modification of their behavior, and we denote this set by $\Omega_r(\xi)$. The larger is ξ, the smaller is the size of this set.

We let $\tilde{\Omega}_r(\xi)$ denote the compliment of this set, and it is only within $\tilde{\Omega}_r(\xi)$ that the size of φ_C can have any effect on household behavior. Any household $\omega \in \tilde{\Omega}_r(\xi)$ has effectively two choices. One is to ignore the program and behave as before. In this case the continuation value of the household remains at $V_{t+1}(\omega)$. If the household alters its investment decisions at time t, by the convexity of the problem, it will choose to change its investment so as to just meet the constraint. If it does so, then the quality level of the child in period $t+1$ (in the absence of any unanticipated productivity shocks) will be $k_{t+1} = \xi$. This child quality level will earn a nonlabor income level next period of $I_{t+1} + \varphi_C$. Then next period's state values for the household will include $w_{1,t+1}$ and $w_{2,t+1}$, which are not affected by the program, $\tilde{I}_{t+1} = I_{t+1} + \varphi$, and child quality of ξ. Let the value of the household's problem as of time $t+1$ with these state variables be given by $\tilde{V}_{t+1}(\omega, \xi, \varphi_C)$. Then it must be the case that $\tilde{V}_{t+1}(\omega, \xi, \varphi_C) > V_{t+1}(\omega)$, since the nonlabor income level is higher in period $t+1$ as is the child quality level than it is without the transfer. Therefore the household's only loss from altering its planned investments in the child in period t from the program associated with ω is its period t loss in utility. Given the household's resources in period t, altering their allocations to achieve the transfer of φ_C in period $t+1$ will result in a lower level of utility in period t, $\tilde{u}_t(\omega, \xi)$. For a household in the set $\tilde{\Omega}_r(\xi)$, the household will increase child cognitive ability in period $t+1$ above its baseline level when

$$\tilde{u}_t(\omega, \xi) + \beta \tilde{V}_{t+1}(\omega, \xi, \varphi_C) \geq V_t(\omega). \tag{12}$$

Then it is clear that the planner faces a number of trade-offs when setting the level ξ. If the level is set too low, then virtually all households qualify for the award and no changes in cognitive outcomes occur. As ξ increases, fewer households will "automatically" qualify for the transfer (i.e., the size of the set $\tilde{\Omega}_r(\xi)$ is increasing in ξ), which means that the policy has the potential to improve child outcomes in more households. As the size of ξ increases, the improvement in child quality in the aggregate will be larger for those who adjust their behavior to receive φ_C. On the other hand, the set of households within $\tilde{\Omega}_r(\xi)$ for whom (12) is satisfied will decrease.

Given the extensive amount of heterogeneity in the model, it is not possible to derive readily interpretable first order conditions for both ξ and φ_C to solve the planner's problem. We can write the cost minimization problem for the conditional transfer problem as follows:

$$\min_{\varphi_C, \xi} \#(\Omega_r(\xi) \cup \tilde{\Omega}_r(\xi, \varphi_C))\varphi_C \tag{13}$$

$$\text{subject to } Z^* = \frac{\int k_{t+1}(s_t, \xi, \varphi_C,) dF_{S_t}(s_t)}{\int k_{t+1}(s_t) dF_{S_t}(s_t)}, \tag{14}$$

where $\tilde{\Omega}_r(\xi, \varphi_C)$ denotes the set of ω who will not automatically qualify for the award and for whom (12) is satisfied. Therefore the population receiving the reward φ_C consists of those households whose baseline level of cognitive ability already meets the threshold ξ ($\Omega_r(\xi)$) and those households whose behavior does not meet the threshold but whose behavior can be altered by the reward φ_C ($\tilde{\Omega}_r(\xi, \varphi_C)$).

This discussion has centered on the ρ_1 standard, which involves only the level of child cognitive ability in period $t+1$, and not its growth level. The case of the ρ_2 is completely analogous, with just a different set of households and actions impacted.

4 Results

In this section we examine the simulation results for the three different transfer program types using the DFW point estimates. Given the sampling design of the data used in the estimation exercise, our estimates can be considered to be nationally representative of all one-child "intact" households. For the unrestricted and restricted transfer case, we examine a number of characteristics of the program by the size of the transfer. The transfer levels we consider in detail are 100, 250, and 500 dollars per week. The characteristics we consider are:

1. The impact on average child quality by the age of the child.
2. The effect of the transfer on average child expenditures, household consumption, parental labor supply, and time spent with the child.
3. The distribution of expenditures across households for the different transfer levels.
4. The distribution of child quality/cognitive ability gains across households for the different transfer levels.

For the case of conditional transfers, we consider the conditional transfer take-up rate by target growth level and reward size. The take-up rate consists of both households for which the criterion would have been reached even under baseline (this is the case of pure rents) and households which adjusted their behavior to satisfy the target (i.e. "compliers" with the program goals). We decompose the take-up rates into these two classes. We also present details of the per capita cost associated with these various combinations of target growth levels and reward sizes.

In the final part of this section, we compare the cost effectiveness of the various transfer policies in the manner outlined in Section 3.

4.1 Unrestricted Transfers

Table 3 presents results for the three levels of unrestricted transfers, \$100, \$250 and \$500. The time unit is weekly, hence a \$100 weekly transfer is equal to a \$5,200 annual transfer. The transfer is given to all households. The transfer is given at age 10; below we discuss how the results vary by the age at which the transfer is given. All figures in the table are average levels across the full population. Below, we discuss the distribution of child quality changes within the population.

Comparing the baseline levels to those with a $100 transfer we see that the average gain in child quality or child cognitive ability (11) is 4.19 / 4.15 = 1.01, about a 1 percent gain from baseline. With a weekly $250 unrestricted transfer, average child quality increases to 4.26, about a 2.7 percent gain from baseline. Finally, a transfer of $500 produces about a 5 percent gain.

The remaining rows in Table 3 show how various household behaviors are affected by the transfer. Turning first to the immediate effect of the transfers on household consumption and child investments, we see that the household consumes part of the transfer and spends part of the transfer on child goods. How much of the transfer is spent on household consumption rather than child goods expenditures depends on the household's preferences for consumption relative to child quality (which is allowed to be heterogeneous in the population as we saw above) and the productivity of child goods expenditures in producing child quality. A particular household may highly value child ability, but at early ages expenditures on child goods are not particularly productive, and thus might instead divert the transfer to household consumption and simultaneously reduce the parents labor supply, spending some of the time saved in child investment. Clearly, the substitution patterns can be quite complex.

This reasoning is confirmed by the results displayed in Table 3. The income transfer does indeed have a large effect on parental labor supply. The time which is taken from labor market activity is spent on time investments with the child and on parental leisure. Both parents spend more time in child investment, in both active and passive activities. The reduction in earned labor income in the household is therefore not necessarily a negative from the point of view of child development, since we have found that in most periods time expenditures of whatever type have a larger impact on cognitive improvements than do money expenditures. However, as we see with the increase in parental leisure time, parents also consume a large portion of the transfer through increased leisure time. The effect of the transfer on time allocation within the household depends on heterogeneous preferences for leisure (for both mothers and fathers), relative wage offers the parents receive (which form the price of foregone working hours), and the productivity of parental time with children, which varies by the age of the child.

Next we explore how the timing of the transfer affects its efficacy. Figure 3 displays the average gain (11) for various levels of unrestricted transfers provided at various child ages. As in all of this analysis, the level of child cognitive ability is given as the next period's output: the transfer is given at age t producing a new level at k_{t+1}. In this Figure, the horizontal axis measures the age t at which the transfer is given, and the horizontal axis measures the gain in average child quality at $t+1$. From the figure we see that for any given level of the transfer, the average gain in child quality is declining with the age at which the transfer is given. Though the model estimates indicate that child investment goods become more productive in child age, we have also seen that there is less "room for improvement" in cognitive ability as the child approaches the end of the development period (as evidenced by the increasing impact of the previous period's ability level). This stasis implies any type of intervention will have more limited impacts at the later stages of the process.

Figure 5 examines the distribution of effects from various levels of unrestricted transfers. Table 3 displays the average level. The figure makes clear that at least for some transfers there is a substantial mass of households with smaller or larger gains than the average. For the relatively small \$100 transfer, most of the gains are centered around the mean, that is, most households experience modest gains from the transfer. For the larger transfers, the gains are more dispersed. For the \$500 transfer, the average gain at the 10th percentile is 1.025, and the average gain at the 90th percentile is 1.083. The difference in household responses is due to the sources of heterogeneity in the model: wage offers, baseline non-labor income, and household preferences. Households with low levels of non-labor income and those with low wage offers respond the most to the transfer. Other households, who already have high levels of household non-labor income or have relatively low preferences for child quality, react relatively little to the program.

4.2 Restricted Transfers

We next turn to the restricted transfer case in which each of the households receive a transfer of φ_R and are required to spend at least φ_R on child goods. As discussed above, some households are unconstrained by this transfer since their baseline expenditure already exceeds φ_R, and for these households the transfer becomes an unrestricted transfer. While for other households with baseline expenditures on children below φ_R, the restriction is binding, as the transfer must be used to fund child goods expenditures rather than be used, at least in part, for household consumption.[12]

Table 4 is the restricted transfer counterpart to Table 3 for the unrestricted transfer case. As in the unrestricted case, the transfer is given to households with a child aged 10. Contrasting the results between the unrestricted and restricted case, we see a larger gain in average child quality through the restricted transfer. A weekly restricted transfer of \$100 results in about a 1.9 percent gain compared to a 1 percent gain for the unrestricted transfer. A weekly restricted transfer of \$500 results in nearly an 11 percent gain in average child quality from baseline compared to about a 5 percent gain for an equal transfer in the unrestricted case.

There are decidedly different household responses to the restricted vs. unrestricted transfers as well. The clearest distinction is in terms of the household allocation of income. With a restricted transfer of \$500, child goods expenditures more than double, a much larger increase than we saw in the unrestricted transfer case in Table 3. As the restricted transfer amount increases, the transfer is increasingly binding for more households, and average child good expenditures are increasing rapidly while household consumption is

[12]The restriction is really only binding for those who with an unrestricted transfer of φ_R would choose to spend less than φ_R on the child. Thus only certain households with a baseline expenditures less than φ_R will be constrained; typically those whose planned baseline expenditures are far below φ_R.

increasing only slightly (as compared to the unrestricted case).

In addition to goods allocation, the overall time allocation response to the restricted transfer is considerably less than for the unrestricted transfer. This is because the restriction on how the income is spent by the household implies that the transfer cannot be optimally (in terms of household utility) allocated across household consumption and child expenditures. The distortion created by the restriction keeps the marginal utility of consumption (and therefore of labor income) higher in the restricted transfer case than in the unrestricted case. The fall in labor hours by the parents is therefore much less in the restricted case than in the unrestricted, as seen in Table 4. As a consequence, the increase in time allocated to children is also less in the restricted transfer case. The time allocation response is nearly flat as the restricted transfer level increases, as compared to the unrestricted transfer case. Transfers that must be spent on child goods do little to change the household's marginal utility of consumption, and hence have little affect on the household's labor supply decisions.

The reason the restricted transfer increases average child quality more than the unrestricted transfer is because the much larger increase in child good expenditures dictated by the restricted transfer outweighs the smaller increase in parental time. While a first order effect of income transfer programs is manifested in the changes in household good allocations, we also see substantial effects of these transfers on time allocation, which can directly affect child outcomes. A perhaps counter-intuitive finding is that an unrestricted transfer is more beneficial in inducing higher parental time investments in children because the unrestricted income can more readily replace parental labor income.

Figure 6 displays results for different levels of the restricted transfer and for different child ages at which the transfer is given. The figure shows that the effect of the transfer is declining in the child's age. Total household income, including both non-labor income and labor income, is increasing as the child ages, hence fewer households have child expenditure levels below the restricted transfer amount. For an increasing share of households, the restricted transfer is essentially an unrestricted transfer, which has a lower overall effect on child quality.

Figure 8 presents the distribution of average quality changes in response to \$100, \$250, and \$500 restricted transfers across the population of households with 10 year old children. While the mean of the distribution is shifted right compared to the unrestricted case at all transfer levels (Figure 5), also notable is that there is a much larger right tail reflecting a considerable mass of households which have quite high gains in average child quality as result of the transfer. These households are constrained households whose baseline level of child expenditures is below the restricted transfer level.

Figure 9 provides an explicit comparison between unrestricted and restricted transfers. At any transfer level, the restricted transfer increases the average gain in child quality significantly more than the unrestricted transfer. As the transfer level increases, the relative advantage of the restricted transfer is larger as the household consumes a higher proportion of the unrestricted transfer.

4.3 Conditional Transfers

Next we turn to results using conditional transfers. In each of these cases, we consider a "value-added" type criterion (ρ_2) provided at age 10 to each household. Unlike the unrestricted and restricted transfer cases, there are two policy variables to consider here: the performance target or threshold, ξ, and the level of reward incentive, φ_C.

Table 5 presents the "take-up" rates for the conditional transfer program. This is the fraction of the population who either already meets the threshold target ξ at baseline (and for whom the reward is a pure rent) and the fraction of the households whose baseline behavior would not satisfy the performance criterion but who would be willing to meet the threshold given the reward incentive φ_C The columns of Table 5 measure the reward incentive and the rows measure the performance target. For completeness, the first entry (target of $\xi = 0$ and reward of $\varphi_C = 0$) corresponds to the baseline, and the take-up rate is by definition 1. As we move down this column (fixing the reward at the baseline of $\varphi_C = 0$), we see what fraction of the population already satisfies the target at baseline. A target of 1 is where the child quality level stays at least at the same level as before ($k_{t+1}/k_t \geq 1$) rather than depreciate at least to some degree. 75 percent of the population meets this threshold. About 51 percent of the population achieves a 10 percent growth in their level of child quality ($k_{t+1}/k_t \geq 1.1$) at baseline. Moving across the columns provides the incentive effect of increasing the one-time reward payment in period $t + 1$ (age 11 in this simulation). We see that with a \$200 weekly equivalent reward incentive, 93.6 percent of the population would meet the threshold of 1.1, compared to only 51 percent at the \$0 reward.

Table 6 displays the fraction of the population that "complies" with the program. The fraction of compliers is the fraction of households who do not meet the target in the baseline but for whom a non-zero reward $\varphi_C > 0$ induces the household to meet the specified target. The compliers can be calculated from Table 5 by subtracting the take-up rate for the non-zero reward columns from the first column for a zero reward $\varphi_C = 0$ (e.g. the fraction of compliers for a target of 1 is 0.83 - 0.75 = 0.08.). The fraction of compliers is relatively low for the low targets because most of the population here would already meet the target at baseline. As would be expected, for a given performance target, the fraction of compliers increases with an increase in the reward.

Table 7 provides the per capita cost of the various combinations of conditional transfer policy parameters. The per capita cost is simply the fraction of the population who would meet the threshold (from table 5) multiplied by the reward level φ_C. In the top row then, with a target threshold of 0, the per capita cost is equal to the reward φ_C since everyone in the population receives the reward (take-up is 1). As the threshold target increases as we move down the columns, the per capita cost declines as the take-up rate declines.

Figure 10 provides the average gain in child quality by the individual household child quality threshold (ξ) for three different levels of reward $\varphi_C = 20, 50, 100, 200$. At low performance targets (ξ small), the difference across reward levels in overall average gain is also relatively small. But as the target level increases, there are much larger differences in the effectiveness of increasing the reward level. At low performance targets, there are few "complier" households who would be unwilling to meet the target at a low reward but who would be willing to meet the target at a high reward. For low performance targets, higher rewards are largely a pure rent for most households. However, as the performance target increases, the fraction of compliers increase as the increasing level of the reward has a greater incentive effect. The larger share of compliers increases the average gain in the cognitive ability level for the total population.

Finally, Figure 11 compares all three types of transfers (unrestricted, restricted, and conditional). For the conditional transfers, we calculate the minimum cost combination of performance targets and rewards to achieve a given average gain (from 14). As the figure shows, for any given gain in average child quality, the policies are clearly ordered in terms of cost effectiveness: conditional transfers are lowest cost, followed by restricted transfers, followed by unrestricted transfers.

The advantage of conditional transfers over the other types of transfers is that the household is allowed to optimally choose the combination of inputs (parental time and child goods expenditures) to meet the performance target. While the household is not child quality maximizing (i.e., it is not choosing the combination of child inputs to maximize the level of child quality each period), and takes into account the utility costs of the input choices (for example the cost of foregone parental leisure in increasing time with children), the household can use the production technology to optimally select the inputs. In contrast, the next best (in terms of cost minimization) policy, restricted transfers, affects child quality primarily through increasing child goods expenditures. This type of transfer distorts the input mix toward child goods expenditures and away from parental time. When the household is left alone to make the decision of which inputs to use, as in the case of the conditional transfer, this distortion is eliminated. Unrestricted transfers suffer from the same type of distortion as the policy distorts the input mix toward goods. An alternative, but infeasible policy, of endowing the parents with more time, would have an opposing kind of distortion of causing the household to sub-optimally increase its input mix toward time with children and away from goods.

5 Conclusion

We have used the model developed and estimated in DFW to examine in some detail the impacts of various types of cash transfer policies on child cognitive outcomes. The model itself is quite stylized, but does incorporate heterogeneity in household preferences, a dynamic child cognitive ability production technology, and a large set of inputs into the production process, which the household chooses given a dynamically evolving resource constraint.

Most social welfare policies, no matter what their goal, involve the transfers of cash to target populations. This is the form of transfers we have considered here. In particular, we looked at cash transfers to intact households with one child that took three forms. First, we considered an unrestricted transfer of φ_U dollars (per week) to all households with children of age t. Second, we considered the transfer of φ_R dollars to all households with children of age t, where the household was restricted in the use of the transfer to the purchase of child investment goods. These might take the form of subsidie for tutors, for example. Finally, we considered the impact of conditional cash transfers, where a household with a child of age t was offered a transfer amount of φ_C if the child's ability at $t+1$ satisfied a certain performance criterion. If it did, the household received the reward of ξ as a weekly transfer amount throughout year $t+1$.

We found that the conditional cash transfer program was considerably more cost-effective than the restricted or unrestricted transfer case. Under the conditional transfer system, some households who would not qualify under baseline must efficiently (in the sense or input utilization) adjust their behavior to (exactly) satisfy the performance criterion and earn the reward ξ. Of course, even among those who qualify for the reward given their baseline behavior will spend some of that "rent" on child quality enhancement in the following period when the reward is received. Thus our case for conditional cash transfers might even be stronger if we looked at longer-run outcomes in the development process.

We have abstracted from important implementation issues and associated costs in our discussion. The unrestricted transfer case, which we found to be the most inefficient, is the easiest one to implement. In the case of restricted transfers, there are costs associated with monitoring the expenditures of households on child investment goods, or costs involved in making in-kind transfers in the amount φ_R to all households. The main costs of the conditional transfer program are in terms of the measurement of cognitive ability and validation of the household having attained the performance threshold. These costs are not trivial, and involve tricky measurement issues from which we have abstracted. For example, if the performance target is set in terms of the change of a measured test score, the household faces uncertainty as to whether its investments will be adequately represented by the scores that the child actually attains. If the household is sufficiently risk averse, the degree of randomness in test scores (conditional on "true" child ability) may make a conditional transfer program considerably less effective than what has been found here.

The main limitation of our results is the restricted environment in which the analysis takes place. The model at present only considers one-child households[13] and short-run transfer programs. Most importantly, other agents in the development process are not considered. Our current research seeks to remedy this problem. In particular, we add school quality to the production technology, and recognize that formal school attendance is the main reason we see such a drop-off in parental time investments around age 6 (and the

[13]In DFW, we extend the model to the case of two-child households and obtain estimates of the more complex production technology for that case. However, performing counterfactual policy analysis is considerably more involved in the two-child case, which is why we have limited our attention to one-child households in this paper.

corresponding implication that the value of parental time investments declines precipitously at this age). Second, we add the child's own time in self-investment to the production technology, and extend the structure of the model to include parental preferences and child preferences separately. In this manner, the counterfactual policy exercise conducted here can be considerably expanded to include incentives extended to schools and children. Deciding not only the amounts of rewards but to whom they should be directed is extremely important, as the results of Behrman et al. (2011) have demonstrated in the case of the ALI social experiment. Our hope is that by combining estimates from models such as the one we have examined and the results of these types of field experiments, more cost-effective policies to promote child development can be found.

References

[1] Aber L, and Rawlings L. (2011). "North South Knowledge Sharing on Incentive-based Conditional Cash Transfer Programs." Social Protection and Labor Discussion Paper 1101, World Bank, Washington, D.C.

[2] Attanasio, O., Battistin, E., Fitzsimons, E., Mesnard, A., and Vera-Hernandez, M.(2005). "How Effective are Conditional Cash Transfers?: Evidence from Colombia." *Institute for Fiscal Studies Briefing* Note 54.

[3] Bernal, R. (2008). "The Effect of Maternal Employment and Child Care on Children's Cognitive Development." *International Economic Review* 49: 1173-1209.

[4] Behrman, J. R., Sengupta, P., and Todd, P. (2005). "Progressing through PROGRESA: An Impact Assessment of a School Subsidy Experiment in Rural Mexico". *Economic Development and Cultural Change* 54: 237-275.

[5] Behrman, J.B., Parker S. Todd, P and Wolpin K. (2011). "Aligning Learning Incentives of Students and Teachers: Results from a Social Experiment in Mexican High Schools." Mimeo, Department of Economics, University of Pennsylvania

[6] Bourguignon, F., Ferreira, F. H. G., and Leite, P. G. (2003). "Conditional Cash Transfers, Schooling, and Child Labor: Micro-simulating Brazil's Bolsa Escola Program." *The World Bank Economic Review* 17: 229-254.

[7] Cardoso, E. and Souza, A. P. (2004). "The Impact of Cash Transfers on Child Labor and School Attendance in Brazil." Working Paper 04-W07, Vanderbilt University Department of Economics.

[8] Cunha, F. and J. J. Heckman. (2008). "Formulating and Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Journal of Human Resources* 43: 738-778.

[9] Cunha, F. and J. J. Heckman, and S. Schennan.(2010). “Estimating the Technology of Cognitive and Non-Cognitive Skill Formation.” *Econometrica* 78: 883-931

[10] Del Boca D, Flinn C. and M. Wiswall (2011). “Household Choices and Child Development.” Carlo Alberto Notebook 149.

[11] Dubois, P., de Janvry, A., and Sadoulet, E. (2008). “Effects on School Enrollment and Performance of a Conditional Transfers Program in Mexico.” WP Series, Department of Agricultural and Resource Economics, UC Berkeley

[12] Fernald L.C., Gertler P.J., and Neufeld L.M. (2008). “Role of Cash in Conditional Cash Transfer Programmes for Child Health, Growth, and Development: An Analysis of Mexico’s Oportunidades.” *Lancet* (9615): 828-37

[13] Fernald LC, Hidrobo M. (2011). “Effect of Ecuador’s Cash Transfer Program on Child Development in Infants and Toddlers: A Randomized Effectiveness Trial.” *Social Science and Medecine* 72: 1437-46.

[14] Heckman, J. J. and Masterov, D. V., (2007). “The Productivity Argument for Investing in Young Children.” IZA Discussion Papers 2725, Institute for the Study of Labor.

[15] Heckman, J., R. Pinto, A. Shaikh, and A. Yavitz (2011). “Inference with Imperfect Randomization: The Case of the Perry School Program.” NBER Working Paper 16935.

[16] Martinelli, C. and Parker, S. W. (2003). ”Should Transfers to Poor Families be Conditional on School Attendance?: A Household Bargaining Perspective.” *International Economic Review* 44: 523-544.

[17] Morris P., Aber L.,Wolf S. Berg J. ”Using Incentives to Change How Teenagers Spend Their Time,The Effects of New York City’s Conditional Cash Transfer Program” MDRC 2012

[18] Paxton C. and Shady N.(2007). “Cognitive Development Among Young Children in Ecuador: The Roles of Wealth, Health and Parenting,” *Journal of Human Resources* 42: 49-84, .

[19] Rawlings, L. B. and Rubio, G. M. (2005). “Evaluating the Impact of Conditional Cash Transfer Programs.” *The World Bank Research Observer*, 20: 29-55.

[20] Schady, N. and Araujo, M. C. (2006). “Cash Transfers, Conditions, School Enrollment, and Child Work: Evidence from a Randomized Experiment in Ecuador.” *World Bank Policy* Research Working Paper, 3930: 1-30

[21] Skofiuas, E. and Parker, S. W. (2001). "Conditional Cash Transfers and Their Impact on Child Work and Schooling: Evidence from the PROGRESA program in Mexico." *Economia* 2: 45-96

[22] Todd, P. E. and Wolpin, K. I. (2006). "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility." *American Economic Review* 96: 1384-1417.

Figure 1: Estimated Child Development Parameters by Child Age (1 Child Model)

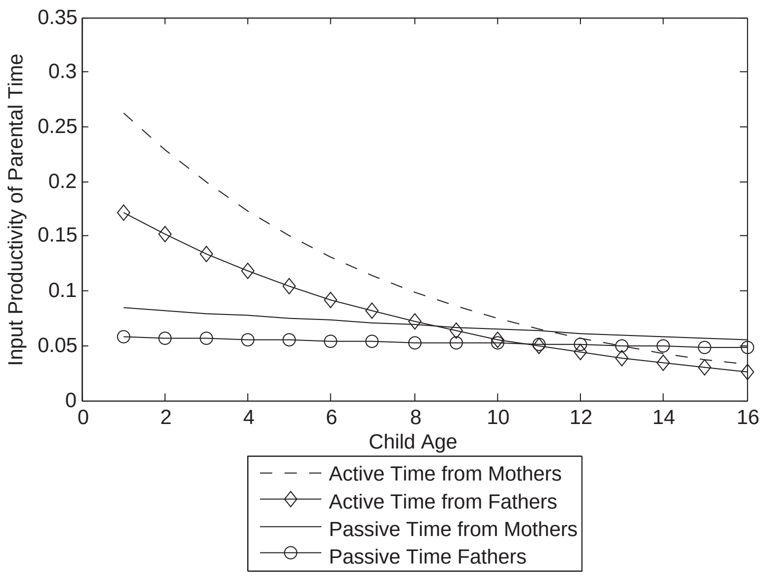

Notes: This graph shows estimated parameters by child age (see DFW 2012).

Figure 2: Estimated Child Development Parameters by Child Age (1 Child Model)

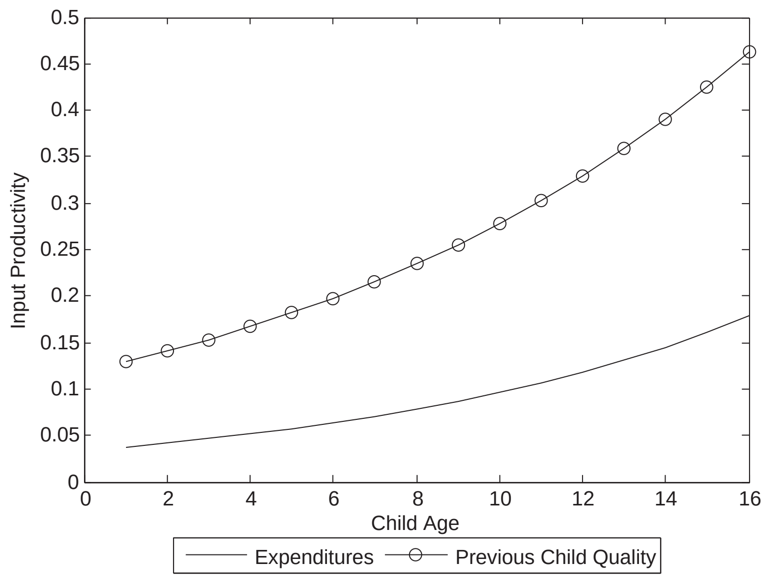

Notes: This graphs estimated parameters by child age (see DFW 2012).

Figure 3: Gain in Average Child Quality from an Unrestricted Income Transfer (by Age of Transfer)

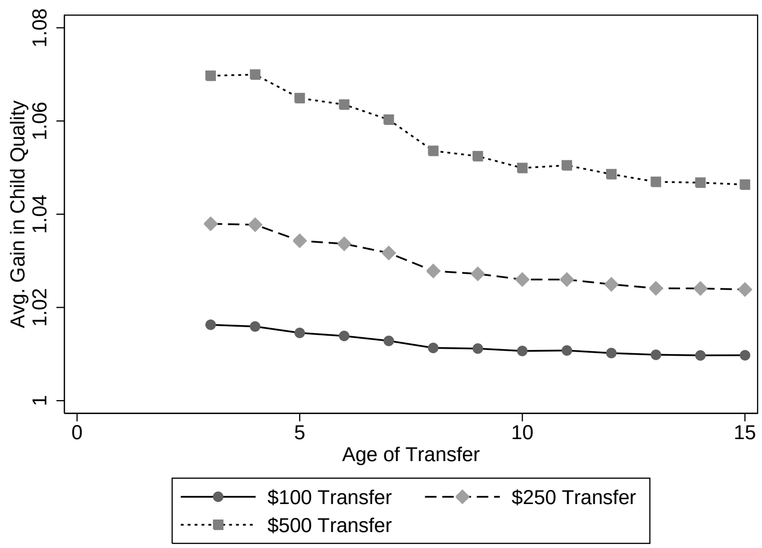

Notes: An unrestricted transfer is an increase in the household's non-labor income, which can be used for any purpose (household consumption or child expenditures). The horizontal axis measures the age at which the transfer is provided (age t). The vertical axis measures the gain in average child quality from baseline as a result of the transfer: $\frac{\int k_{t+1}(s_t,\varphi(\zeta))dF_{S_t}(s_t)}{\int k_{t+1}(s_t)dF_{S_t}(s_t)}$, where $\int k_{t+1}(s_t)dF_{S_t}(s_t)$ is the baseline level of average child quality and $\int k_{t+1}(s_t,\varphi(\zeta))dF_{S_t}(s_t)$ is the average level with the transfer. Inputs provided at age t produce child quality at age $t+1$. All figures are produced using simulation from the parameter estimates of the one child model.

Figure 4: Distribution of Child Expenditures following Unrestricted Transfers

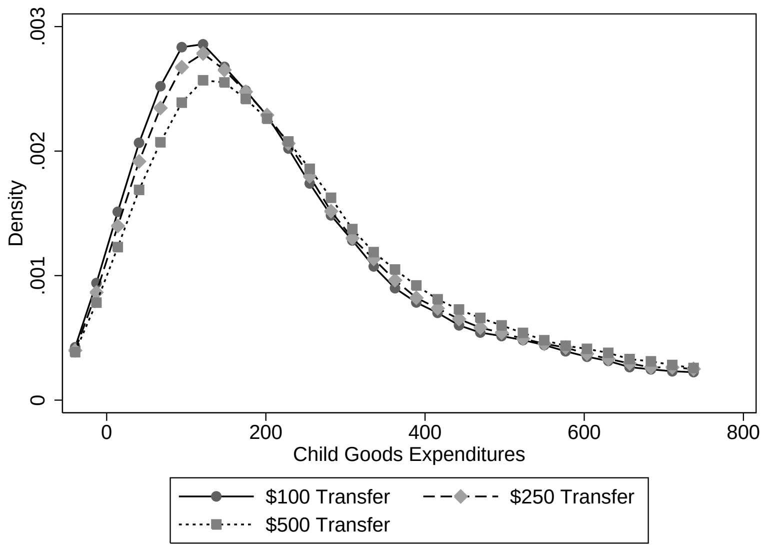

Notes: An unrestricted transfer is an increase in the household's non-labor income, which can be used for any purpose (household consumption or child expenditures). The figure plots a smoothed density of expenditures on child goods. Transfer is provided at age 10 and the figure shows expenditures at age 10 following this transfer. All figures are produced using simulation from the parameter estimates of the one child model.

Figure 5: Distribution of Gain in Child Quality from an Unrestricted Transfer

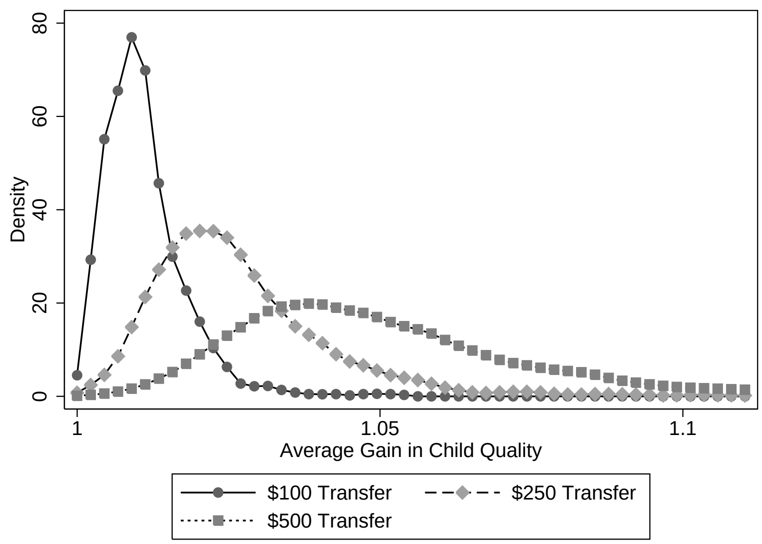

Notes: An unrestricted transfer is an increase in the household's non-labor income, which can be used for any purpose (household consumption or child expenditures). The figure plots a smoothed density of the gain in child quality from baseline as a result of the transfer. Inputs provided at age t produce child quality at age $t+1$. Transfer is provided at age 10 ($t = 10$) and the figure shows child quality for age 11 (k_{11}). All figures are produced using simulation from the parameter estimates of the one child model.

Figure 6: Gain in Average Child Quality from a Restricted Income Transfers (by Age of Transfer)

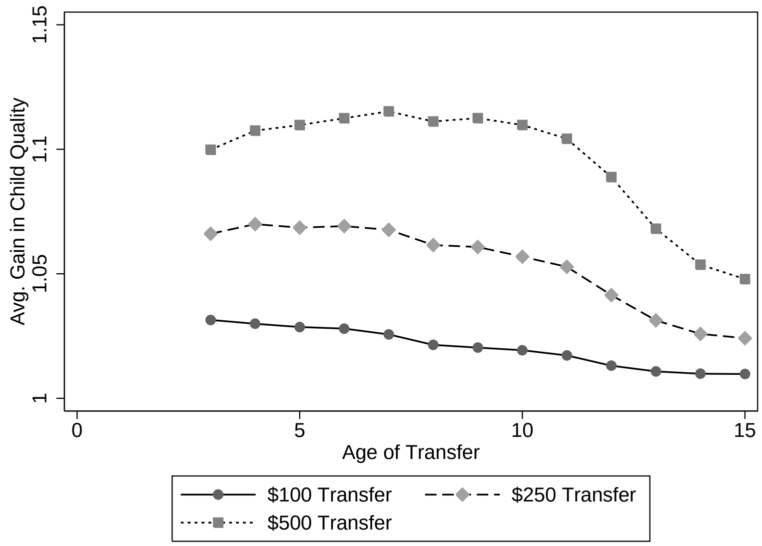

Notes: A restricted transfer is a transfer of φ_R in non-labor income with a restriction that the household spends at least φ_R on child goods. The horizontal axis measures the age at which the transfer is provided (age t). The vertical axis measures the gain in average child quality from baseline as a result of the transfer: $\frac{\int k_{t+1}(s_t,\varphi(\zeta))dF_{S_t}(s_t)}{\int k_{t+1}(s_t)dF_{S_t}(s_t)}$, where $\int k_{t+1}(s_t)dF_{S_t}(s_t)$ is the baseline level of average child quality and $\int k_{t+1}(s_t,\varphi(\zeta))dF_{S_t}(s_t)$ is the average level with the transfer. Inputs provided at age t produce child quality at age $t+1$. All figures are produced using simulation from the parameter estimates of the one child model.

Figure 7: Distribution of Child Expenditures following Restricted Transfers

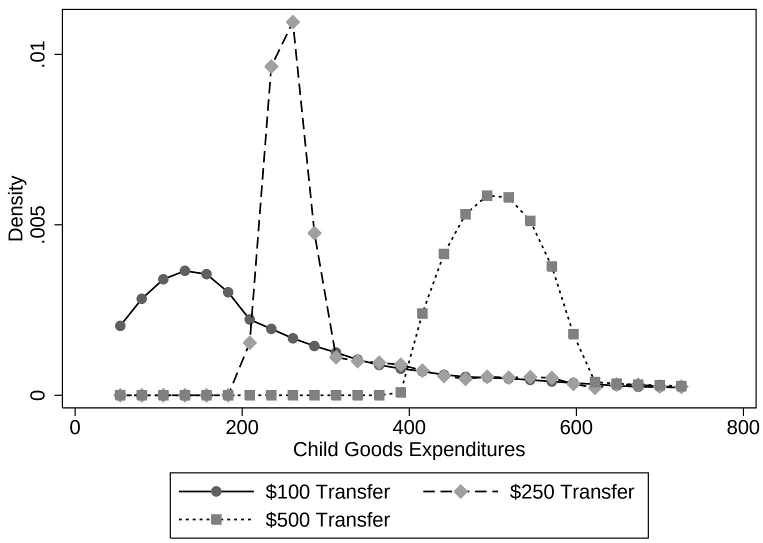

Notes: A restricted transfer is a transfer of φ_R in non-labor income with a restriction that the household spends at least φ_R on child goods. The figure plots a smoothed density of expenditures on child goods. Transfer is provided at aged 10 and the figure shows expenditures at age 10 following this transfer. All figures are produced using simulation from the parameter estimates of the one child model.

Figure 8: Distribution of Gains in Child Quality from a Restricted Transfer

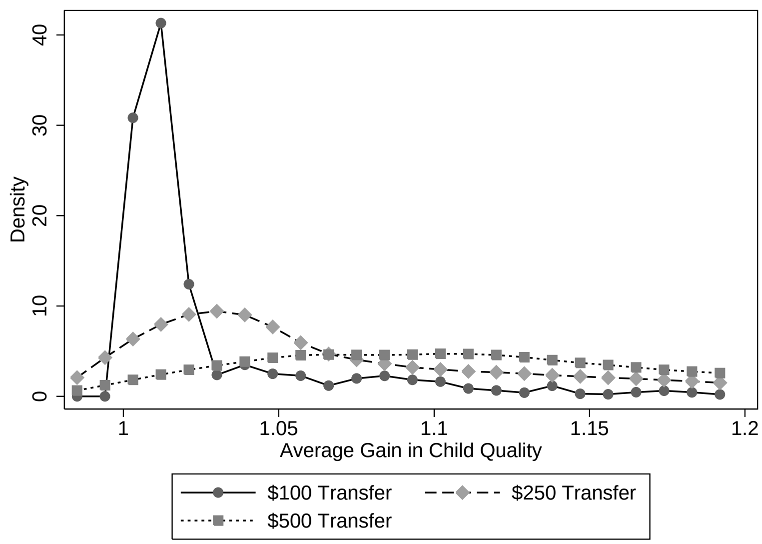

Notes: A restricted transfer is a transfer of φ_R in non-labor income with a restriction that the household spends at least φ_R on child goods. The figure plots a smoothed density of the gain in child quality from baseline as a result of the transfer. Inputs provided at age t produce child quality at age $t+1$. Transfer is provided at age 10 ($t = 10$) and the figure shows child quality for age 11 (k_{11}). All figures are produced using simulation from the parameter estimates of the one child model.

Figure 9: Comparison of Unrestricted vs. Restricted Transfer

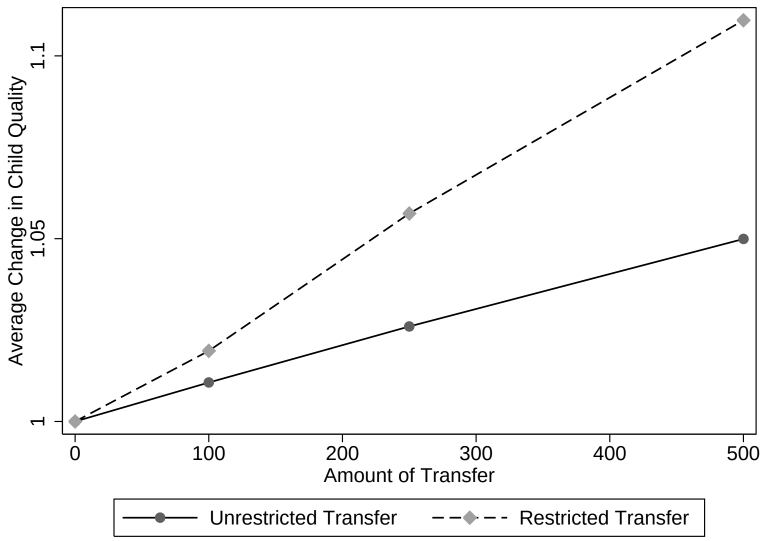

Notes: An unrestricted transfer is an increase in the household's non-labor income, which can be used for any purpose (household consumption or child expenditures). A restricted transfer is a transfer of φ_R in non-labor income with a restriction that the household spends at least φ_R on child goods. Transfer is provided at age 10 ($t = 10$). Child quality is the the level of child quality produced at the end of age 10 and the initial value for age 11 (k_{11}). All figures are produced using simulation from the parameter estimates of the one child model.

Figure 10: Average Gain in Child Quality of Conditional Transfer Program

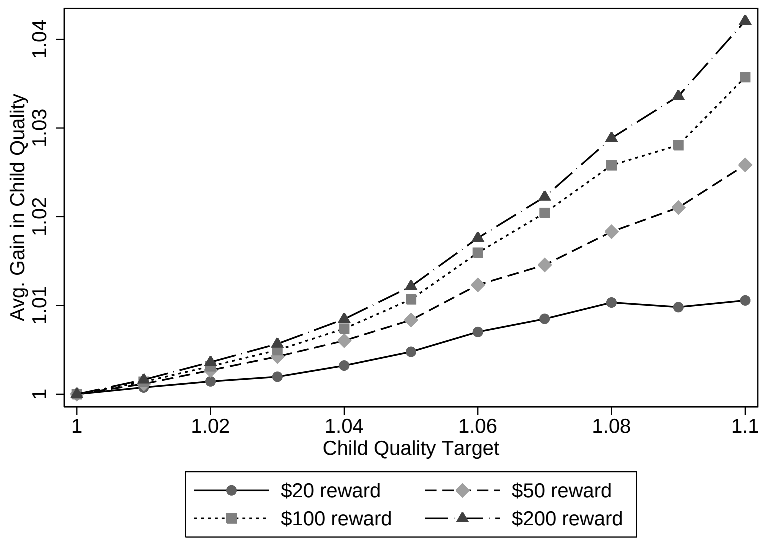

Notes: A conditional transfer program transfers φ_C dollars to a household in period $t+1$ if the household meets the target of increasing child quality by $k_{t+1}/k_t \geq \xi$. In this Figure we fix age t to be age 10.

Figure 11: Comparison of Unrestricted, Restricted, and Conditional Transfers

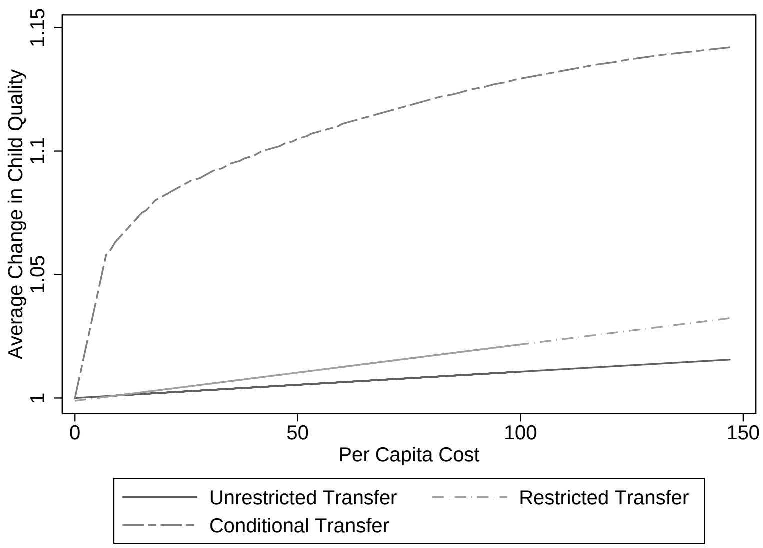

Notes: An unrestricted transfer is an increase in the household's non-labor income, which can be used for any purpose (household consumption or child expenditures). A restricted transfer is a transfer of $X in non-labor income with a restriction that the household spends at least $X on child goods. Transfer is provided at age 10 ($t = 10$). A conditional transfer program transfers φ_C dollars to a household in period $t+1$ if the household meets the target of increasing child quality by $k_{t+1}/k_t \geq \xi$. The conditional policy here is the minimum cost combination of policy target and reward to achieve a given improvement in average child quality. Child quality is the the level of child quality produced at the end of age 10 and the initial value for age 11 (k_{11}). All figures are produced using simulation from the parameter estimates of the one child model.

Table 1: Preference Parameter Moments for One-Child Families

Parameter	Description	Mean	Coef. of Variation
α_1	Mother's Leisure	0.196	0.619
α_2	Father's Leisure	0.194	0.440
α_3	Household Consumption	0.257	0.362
α_4	Child "Quality"	0.353	0.568

Notes: Table reports summary of moments of estimated preference parameters (DFW 2012).

Table 2: Effects of an Increase in the Parents' Wages on Household Outcomes

	Level at Baseline	Increase in Mother's Wage	Increase in Father's Wage
		Elasticity	
Mean Latent Child Quality (Age 16)	13.328	0.093	0.142
Mean Hours Work (Mother)	30.205	1.092	-1.054
Mean Hours Work (Father)	44.316	-0.596	0.609
Mean Active Time w/ Child (Mother)	16.890	-0.445	0.435
Mean Active Time w/ Child (Father)	11.056	0.397	-0.394
Mean Passive Time w/ Child (Mother)	13.364	-0.456	0.446
Mean Passive Time w/ Child (Father)	9.732	0.415	-0.415
Mean Leisure (Mother)	51.541	-0.377	0.358
Mean Leisure (Father)	46.916	0.379	-0.384
Mean Child Expenditures / 1000	0.377	0.437	0.531
Mean Household Consumption / 1000	1.110	0.410	0.559

Notes: This table reports elasticity estimates from an increase in mother's and father's wage offer. Mean Latent Child Quality (Age 16) is the latent value of child quality at the end of age 16 or the start of period $t = 17$, k_{17}.

Table 3: Unrestricted Transfer

	(1) No Transfer	(2) 100 Transfer	(3) 250 Transfer	(4) 500 Transfer
Child Quality (Latent)	4.15	4.19	4.26	4.36
Mother's Labor Hours	33.4	31.2	28.4	24
Father's Labor Hours	42.8	40.6	37.1	31.8
Mother's Active Time	12.5	12.9	13.5	14.3
Father's Active Time	8.29	8.55	8.97	9.61
Mother's Passive Time	10.8	11.2	11.6	12.3
Father's Passive Time	7.7	7.94	8.32	8.92
Mother's Leisure Time	55.3	56.8	58.5	61.4
Father's Leisure Time	53.2	54.9	57.6	61.6
Child Expenditures	261	268	278	298
Household Consumption	1261	1297	1349	1445

Notes: An unrestricted transfer is an increase in the household's non-labor income, which can be used for any purpose (household consumption or child expenditures). Transfer is provided at age 10 ($t = 10$). Child quality is the the level of child quality produced at the end of age 10 and the initial value for age 11 (k_{11}). All figures are produced using simulation from the parameter estimates of the one child model.

Table 4: Restricted Transfer

	(1) No Transfer	(2) 100 Transfer	(3) 250 Transfer	(4) 500 Transfer
Child Quality (Latent)	4.15	4.23	4.38	4.6
Mother's Labor Hours	33.4	31.4	30	29.1
Father's Labor Hours	42.8	41	39.5	38.6
Mother's Active Time	12.5	12.9	13.3	13.6
Father's Active Time	8.29	8.54	8.81	9.02
Mother's Passive Time	10.8	11.1	11.5	11.7
Father's Passive Time	7.7	7.93	8.18	8.38
Mother's Leisure Time	55.3	56.5	57.2	57.6
Father's Leisure Time	53.2	54.5	55.5	56
Child Expenditures	261	279	354	550
Household Consumption	1261	1292	1318	1338

Notes: A restricted transfer is a transfer of φ_R in non-labor income with a restriction that the household spends at least φ_R on child goods. Transfer is provided at age 10 ($t = 10$). Child quality is the the level of child quality produced at the end of age 10 and the initial value for age 11 (k_{11}). All figures are produced using simulation from the parameter estimates of the one child model.

Table 5: Conditional Transfer Take-up Rates

	(1)	(2)	(3)	(4)	(5)
	Transfer Amount:				
	0	20	50	100	200
Target					
0	1	1	1	1	1
1	.75	.83	.884	.93	.96
1.01	.744	.822	.876	.92	.958
1.02	.728	.806	.874	.92	.958
1.05	.674	.768	.844	.904	.954
1.1	.51	.612	.758	.862	.936

Notes: Take-up rate is the fraction of the population who either already meets the threshold target ξ at baseline or the households who would be willing to meet the threshold given the reward incentive φ_C. The columns measure the reward incentive and the rows measure the performance target. For completeness, the first entry (target of $\xi = 0$ and reward of $\varphi_C = 0$) corresponds to the baseline, and the take-up rate is by definition 1. As we move down this column (fixing the reward at the baseline of $\varphi_C = 0$), we see what the baseline fraction of the population that already meets the target. A target of 1 is where the child quality level stays at least at the same level as before ($k_{t+1}/k_t \geq 1$) rather than depreciate at least to some degree. A target of 1.01 is a 1 percent gain in child quality.

Table 6: Conditional Transfer Fraction of Compliers

	(1)	(2)	(3)	(4)
	Transfer Amount:			
	20	50	100	200
Target				
0	0	0	0	0
1	.08	.134	.18	.21
1.01	.078	.132	.176	.214
1.02	.078	.146	.192	.23
1.05	.094	.17	.23	.28
1.1	.102	.248	.352	.426

Notes: The fraction of compliers are the fraction of households who not meet the target in the baseline but for which a non-zero reward $\varphi_C > 0$ induces the household to meet the specified target. The compliers can be calculated from Table 5 by subtracting the take-up rate for the non-zero reward columns from the first column for a zero reward $\varphi_C = 0$ (e.g. the fraction of compliers for a target of 1 is 0.83 - 0.75 = 0.08.).

Table 7: Conditional Transfer Per Capita Costs

	(1)	(2)	(3)	(4)	(5)
	Transfer Amount:				
	0	20	50	100	200
Target					
0	0	20	50	100	200
1	0	16.6	44.2	93	192
1.01	0	16.4	43.8	92	192
1.02	0	16.1	43.7	92	192
1.05	0	15.4	42.2	90.4	191
1.1	0	12.2	37.9	86.2	187

Notes: Per capita cost is the fraction of the population who would meet the threshold (from table 5) multiplied by the reward level φ_C. In the top row then, with a target threshold of 0, the average cost is equal to the reward since everyone in the population receives the reward (take-up is 1). As the threshold target increases moving down the rows, the per capita cost declines as the take-up rate declines.