

IZA DP No. 8393

Transfers to Households with Children and Child Development

Daniela Del Boca
Christopher Flinn
Matthew Wiswall

August 2014

Transfers to Households with Children and Child Development

Daniela Del Boca

University of Torino, CCA and IZA

Christopher Flinn

New York University and CCA

Matthew Wiswall

Arizona State University

Discussion Paper No. 8393

August 2014

IZA

P.O. Box 7240

53072 Bonn

Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Transfers to Households with Children and Child Development^{*}

In this paper we utilize a model of household investments in the development of children to explore the impact of various transfer policies on the distribution of child outcomes. 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 an optimally chosen conditional cash transfer program is the most cost efficient way to attain any given gain in average child quality. We explore several design elements for the conditional cash transfer system and discuss the role of production function uncertainty and measurement error.

JEL Classification: J13, D1

Keywords: child development, time allocation, income transfers,
conditional cash transfers

Corresponding author:

Daniela Del Boca
Department of Economics
University of Turin
Via Po 53
Torino 10124
Italy
E-mail: dani.delboca@unito.it

^{*} We acknowledge support from the National Science Foundation, the Collegio Carlo Alberto ("Parental Choices and Child Outcomes"), the Institute of Human Development and Social Change, and the C.V. Starr Center for Applied Economics. We have benefited from the comments of two anonymous referees, seminar participants at Colorado, and conference participants at the Yale Structural Microeconomics conference (June 2014) and the IFS Workshop on Risk and Family Economics (June 2014). We are responsible for all errors, interpretations, and omissions.

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 parents lack either resources, parental skills, 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 child development, and there exists important research on the optimal timing of investments (early versus late childhood) and the types of skills in which to invest (cognitive versus non-cognitive), there is far less agreement about the specific types of interventions which are likely to be most effective.

In this paper, we utilize a model of household behavior and child development to compare a broad class of transfer-based interventions. We consider three types of transfer programs: (i) an “unrestricted” transfer in which the household receives a lump sum transfer with no restrictions on its use; (ii) a “restricted” transfer of child investment goods or services such as would be embodied in a Perry or Abecedarian type program; and (iii) a “conditional” (cash) transfer, or CCT, in which the transfer is made to the household only after the child’s measured development satisfies some performance criteria.¹ We do not specify a complete social planning problem, but confine our attention to the evaluation of the cost effectiveness of the transfer programs in producing a given change in the distribution of child outcomes. Within each class of policies, we study the issues related to the design of the transfer, including considerations of the age at which to provide the transfer, and in the case of the more complex CCTs, how different performance criteria and reward levels change incentives for child investment.

This paper builds on our earlier work, Del Boca, Flinn, and Wiswall (2014, hereafter DFW), in which we estimate a simple model of investments in child quality with heterogeneous household preferences, resource constraints, and a growth process for the child’s development. A key feature of our approach is that we explicitly model a variety of parental investments in children: both mother’s and father’s time with children and expenditures

¹We analyze only output based CCT programs, where the criteria for receiving the reward is based on the level of child development. Other research explores “input” based CCT programs, where the criteria for receiving the reward is based on the household taking some particular actions. Some other research uses the term “conditional” to refer to transfers conditioned on some observed baseline characteristic of the household such as the baseline level of household income (e.g. a “means tested” transfer policy provided only to poor families). We refer to these policies as “targeted” and discuss this issue below. Our use of the term “conditional” refers to the transfer conditioned on some action of the household or some outcome that is produced by their actions (i.e. next period’s child development level).

on child goods. We jointly model these investments with household labor, leisure, and parental consumption choices, allowing us to examine several margins of household responses to interventions. In using our model to compare policies, we note two key tradeoffs in the design of policies to affect child development. First, policies differ in how much households “consume” the transfer (through purchases of consumption goods or leisure) rather than using the transfer to invest in children (through purchases of child-specific goods or time investments). Second, policies differ in the extent to which they distort the household’s investments in children away from the optimal mix of inputs (goods and time of different agents) that would best affect child outcomes. In our quantitative results, we conclude that CCTs are the most cost effective in improving average child cognitive ability since they have more limited scope for household consumption relative to an unrestricted transfer and distort the child input mix less than restricted transfers of child goods. We discuss a number of additional considerations, such as administrative costs, dynamic “ratchet” effects, and measurement issues, which may be important enough to mitigate this conclusion, but which are not considered in our quantitative analysis.

Heckman and Mosso (2014) provide an extensive review of research on unrestricted transfers to households and the more general issue of the relationship between family income and child development. The strong positive correlation between family income and child development does not necessarily indicate that an unrestricted income transfer would have a correspondingly large effect on children. Several recent papers use various sources of exogenous variation in family income to estimate a positive but modest “reduced-form” effect of family income on measures of child outcomes (Dahl and Lochner 2012; Loken, Mogstad, and Wiswall (2012); Duncan, Morris, and Rodriguez 2011). In general, the reduced-form relationship is not easily interpretable given that for most households income is primarily generated by labor market earnings, and these require substantial time commitments from parents. To the extent that parental time investments are important factors in child development, this tends to decrease the resources devoted to the children. For some households, this channel may dampen or even reverse the assumed positive relationship between income and child development.

“Restricted,” or in-kind, transfers of resources directly to children have been studied extensively but usually by evaluating particular programs. Much of the research in the United States and other developed countries has focused on interventions that provide children with better environments outside of the home. Heckman and Kautz (2013) provide a recent summary of many of these programs. 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.

Conditional cash transfer programs (CCTs) have been increasingly implemented in developing countries since the 1990s. For the most part, these CCT programs have been

“input-based,” with the household being rewarded with some transfer for household behaviors such as sending the child to a health clinic or school. 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. 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 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. (2012), and Todd and Wolpin (2006)). The effects of CCTs on learning and school achievement are less widely studied. 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 CCTs, with some running pilot studies. 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 significantly affect school outcomes for younger children, it substantially improved the achievement of older children (Morris et al 2012).

A common problem faced by the 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. (2014). Here the authors report the results of their social experiment in Mexico, 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-executed, 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.²

The plan of the paper is as follows. In Section 2 we lay out the model structure from DFW (2014). Section 3 provides a formal discussion of the household's responses to various policies. Section 4 compares the three types of transfer systems we consider under a cost minimization criteria. Section 5 contains the results from our simulations and the effectiveness comparisons across the three transfer types. Section 6 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.

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. Although we also 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 one of the best (estimated) model of the child development process currently available. Cunha (2013) and Caucutt and Lochner (2012) consider models of child development with several features not found in our model: borrowing and saving and multiple generations,

²The model we exposit below does not include formal schooling, so school-based incentive systems are not considered. 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.

and in Cunha’s case, the model is solved for a steady state equilibrium. However, these papers consider a more limited number of inputs (only child goods expenditures) and do not model the time allocation of the household including time spent with children or labor supply decisions. While our framework does not allow us to examine the impacts of policy on future generations nor in general equilibrium, our model does allow a rich variety of household responses (both in goods and time allocation) to the policies we consider. We discuss briefly the importance of credit constraints on household responses to transfer policies below.

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 Cunha and Heckman (2008) and Cunha et al. (2010) has been extremely influential in establishing the importance of cognitive and 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 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.³

³The 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.

In each period, the household makes ten choices: the hours of work for each parent: h_{1t} (mother) and h_{2t} (father); time spend in leisure by both parents, l_{1t} and l_{2t} ; 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 $\tau_{2t}(p)$; expenditures on “child” goods, e_t ; and expenditures on market goods consumed by the household, c_t . Household utility in period t is a function of l_{1t} , l_{2t} , c_t , and k_t , the level of the child’s quality at the beginning of period t . We assume a Cobb-Douglas household utility function and restrict the preference parameters to be stable over time:

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

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,

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

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

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 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 The Household's 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. We assume a unitary household utility function that is Cobb-Douglas, so that $u(l_{1,t}, l_{2,t}, c_t, k_t) = \alpha_1 \ln l_{1,t} + \alpha_2 \ln l_{2,t} + \alpha_3 \ln c_t + \alpha_4 \ln k_t$. The value function for the household in period t is then

$$V_t(S_t) = \max_{i_t} \alpha_1 \ln l_{1,t} + \alpha_2 \ln l_{2,t} + \alpha_3 \ln c_t + \alpha_4 \ln k_t + \beta E_t V_{t+1}(S_{t+1}), \quad (3)$$

$$\text{s.t. } T = l_{jt} + h_{jt} + \tau_{jt}(a) + \tau_{jt}(p), \quad j = 1, 2$$

$$c_t + e_t = w_{1t}h_{1t} + w_{2t}h_{2t} + I_t$$

where i_t is the vector of decision variables. The vector of state variables in period t consists of the current level of child quality, the wage offers to the parents, and nonlabor income, so that $S_t = (w_{1t} \ w_{2t} \ I_t \ k_t)$. The discount factor is β ($\in [0, 1)$), 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, $S_1 = (k_1 \ w_{11} \ w_{21} \ I_1)$. To simplify the analysis, for the most part we will be assuming that the household has perfect foresight regarding future values of S . Under our modeling assumptions, there is no loss of generality in doing so except in the case of conditional cash transfer systems, upon which we will comment at the appropriate point.

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 of 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 periods 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 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 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

$$V_{M+1}(S_{M+1}) = \psi \alpha_4 \ln k_{M+1},$$

where ψ is a parameter to be estimated.

2.5 Model Solution

For each household, the model is solved for the optimal parental inputs and labor supply in each development period from $t = 1$ to $t = M$, given by the vectors

$$\{h_{1t}^*, h_{2t}^*, l_{1t}^*, l_{2t}^*, \tau_{1t}(a)^*, \tau_{2t}(a)^*, \tau_{1t}(p)^*, \tau_{2t}(p)^*, e_t^*, c_t^*\}_{t=1}^M.$$

These optimal choices determine the level of the child's ability, denoted $\{k_t^*\}_{t=1}^{M+1}$. As is clear from the nature of the production technology, there are never any corner solutions in the household input choice problem during the investment period.⁴ However, we do allow for corner solutions in labor supply since either or both parents may decide not to participate in the labor market in a given period. The explicit form of the decision rules is provided in the Appendix A.

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. However, in order to enhance the precision of the estimator we use and for intertemporal consistency of the decision rules, we have assumed time-invariant household preferences.

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.

⁴If any factor is set at 0, then child quality will be 0 in all subsequent periods, and household utility diverges to $-\infty$ as $k \rightarrow 0$ whenever $\alpha_4 > 0$.

2.6 Empirical Specification

In taking the model to data, we allow heterogeneity in several dimensions: i) preferences, ii) initial child quality, iii) wage offers for parents, and iv) non-labor income. In terms of preference heterogeneity, we assume that the household preferences for leisure, consumption, and child quality $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ have a (latent) joint Normal distribution that respects the normalizations $\alpha_j \in (0, 1)$ for all j and $\sum_j \alpha_j = 1$.⁵ The child development technology varies by the child’s age and the initial child quality endowment, but is otherwise homogenous.⁶

2.7 Data and Estimation

We utilize data from the Panel Study of Income Dynamics (PSID) and the first two waves of the Child Development Supplements (CDS-I and CD-II). The PSID is a longitudinal study that began in 1968 with a nationally representative sample of about 5,000 American households, with an oversample of black and low-income households. In 1997, the PSID began collecting data on a random sample of the PSID families that had children under the age of 13 in the Child Development Supplement (CDS-I).

Beginning in 1997, children’s time diaries were collected along with detailed assessments of children’s cognitive development. For two days per week (one weekday and either Saturday or Sunday), children (with the assistance of the primary caregiver when the children were very young) filled out a detailed 24 hour time diary in which they recorded all activities during the day and who else (if anyone) participated with the child in these activities. At any point in time, the children recorded the intensity of the participation of the parents in the activities: mothers and fathers could be actively participating or engaged with the child or simply around the child but not actively involved. We refer to the first category of time as “active” time and the second as “passive.” We then utilize four categories of time inputs, active and passive time spent with each of the parents. We construct a weekly measure of each type of child investment time for the mother and father by multiplying the daily hours by 5 for the weekday and 2 for the weekend day (using a Saturday and Sunday report adjustment) and summing the total hours for each category of time.

⁵In particular, we draw from a trivariate normal distribution a vector x , which has mean μ_x and covariance matrix Σ_x . From this we form $\alpha_1 = \exp(x_1)/D$, $\alpha_2 = \exp(x_2)/D$, $\alpha_3 = \exp(x_3)/D$, and $\alpha_4 = 1/D$, where $D \equiv \exp(x_1) + \exp(x_2) + \exp(x_3) + 1$. Thus for all draws, each $\alpha_i > 0$, $i = 1, \dots, 4$, and $\sum \alpha_i = 1$. The “primitives” of the preference specification are the mean μ_x and the covariance matrix Σ_x which generate the population distribution of Cobb-Douglas preference parameters.

⁶In an earlier version of this paper, we estimated a model in which mother’s education affected the productivity of the mother’s time with children, and similarly for the father. We found this model was weakly identified, which we suspect is due to the fact that mother’s education also affects her wage offer and therefore indirectly also determines her optimal time inputs. Perhaps more importantly, the sample size used in the estimation is quite small. We believe that estimation with larger samples would resolve the weak identification problem and lead to production function estimates consistent with our priors.

We measure the child’s skills using the Letter-Word component of the Woodcock Johnson Achievement Test-Revised (Woodcock and Johnson, 1989). This test has the attractive property that it is appropriate for children as young as 3 and as old as 16. We use the raw scores on this exam rather than the age-standardized scores. The test contains 57 items, with the range of possible raw scores being from 0 to 57. We utilize a measurement error model that respects the discrete nature of the test score with a finite ceiling and floor.

We are interested in households in which both biological parents were present in both waves. Most of the variables we use in the model are collected from the primary caregiver of a child and for the head and wife of the household. Therefore, our initial sample selection results in households with children in the CDS who (1) have valid test scores in both the 1997 and the 2002 waves of the CDS, and (2) are sons or daughters of the head of the household. In addition we drop observations with missing information on mother’s or father’s time with the child or missing age or education of either parent. We do not use wage observations if the reported (real) hourly wage is more than \$150 per hour, and do not use an income observation if the reported weekly nonlabor income is greater than \$1,000. Our total sample consists of 105 one-child households. The observed household characteristics include parental variables, such as the education and the ages of the parents when the child was born. For each mother and father in the household we observe: hours worked, (accepted) wages for both parents, and non-labor income.

We estimate the parameters of the model using the Method of Simulated Moments. For each observed household in the data, we stochastically draw from the preference distribution and the measurement error process to simulate the sequence of optimal household choices and measured child cognitive ability. While we do not observe, choices and child outcomes in each period, we use the model structure to “fill-in” the missing data using Monte Carlo integration. With the simulated data set, we then compute the analogous simulated sample characteristics to those determined from the actual data sample, and form a method of moments estimator. The moments we use include the average and standard deviation of test scores at each child age, the average and standard deviation of hours of work for mothers and fathers at each child age, and the average and standard deviation of child investment hours for mothers and fathers at each child age. In addition, we use the average and standard deviation of accepted wages and the correlation in wages across parents. We also include a number of contemporaneous and lagged correlations between the observed labor supply, time with children, child quality, wages, and income.

2.8 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 and parent 2 is the father. The first thing to note is that households, on averages, attach the largest utility weight to their child’s ability. However, the weight attached is less than the sum of the weights attached to each parent’s leisure, and is only about 37 percent greater than the weight attached to household consumption. Although 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 partially 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 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 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.

3 Household Responses to Transfers

In this section we consider the impact on child development of three types of transfer programs. 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.⁷ We compare each of these policies in terms of the cost to achieve some gain in the aggregate latent child quality stock of children when transitioning from age t to age $t + 1$. In the following sections, it will be useful to denote a household type when the child is age t by ω_t , with the distribution of household types in the population given by $F_{\Omega_t}(\omega_t)$. The household's type when the child is aged t includes all of the household state variables $(w_{1,t}, w_{2,t}, I_t, k_t)$ as well as its preference weights $\alpha = (\alpha_1 \alpha_2 \alpha_3 \alpha_4)$. We do not include the production technology coefficients δ_t and TFP, R_t , since these parameters are common to all households with a child of age t .

3.1 Unrestricted Transfers

In the unrestricted transfer policy, indexed P_U , households are transferred some amount with no restriction on how it should be spent. In this policy, household's have their nonlabor income increase at child age t from I_t to $I_t + \phi_U$, where $\phi_U \in (0, \infty)$ is the increase in non-labor income for the policy P_U . Because the transfer is an increase in non-labor income, our behavioral model indicates that 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. As discussed in more detail below, the proportion of the transfer "consumed" by the parents depends on the household's preferences, with households who value their own consumption and leisure relatively more than child quality consuming more of the transfer.

The unrestricted lump sum transfer that we consider here has as a benefit the simplicity of the transfer policy and the lack of enforcement required to implement it. In addition, we are only considering a policy in which all households in the target population receive exactly the same transfer, ϕ_U . It is, of course, possible to make the transfer amount a function of household characteristics, at least the subset of these characteristics observable to the planner. We will discuss this issue in a bit more depth at the end of the next section.

⁷The 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 more 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.

3.2 Restricted Transfers

A restricted transfer policy, indexed by P_R , is the same as an unrestricted one except that the institutional agent requires the household to spend at least the amount $\phi_R \in (0, \infty)$ on the child, so that child expenditures $e_t \geq \phi_R$. 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. This could take the form of an enrichment program directly provided to children as with the Perry or Abcederian programs. 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.⁸ All households given a restricted transfer receive a higher utility level than in the absence of any transfer, though the utility level can be no higher than it would be if the transfer were unrestricted. As is true for the case of unrestricted transfers, we consider only the simplest form of this system in which all targeted households receive the same transfer amount.

Although this type of transfer acts as a constraint on the allocation decisions of some households, for some households the restricted transfer has the same effect on household behavior as an unrestricted transfer. Say that prior to notification of this transfer, the household plans an expenditure of e_t . Then if $e_t \geq \phi_R$, 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 . 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.⁹

3.3 Conditional Cash Transfers

In the conditional cash transfer (CCT) policy, indexed P_C , the household only receives the transfer amount if it qualifies on the basis of some performance criteria. In many CCT 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. In contrast to these input-based transfer programs, we examine output-based conditional transfers that are made on the basis of the child's cognitive ability.

We will develop a fairly general framework within which we can analyze household behavior under a variety of CCT reward and performance structures. We will, however,

⁸This 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 those used in producing child quality.

⁹For 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 distorting relative to an unrestricted transfer. 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 household decisions.

only be considering very specific CCT environments. Although it is conceptually straightforward to extend our framework to more general cases, computationally the task becomes much more difficult. Since our desire is to illuminate general issues in the implementation of CCTs and compare these programs to the unrestricted and restricted transfer programs, we will work with the simplest environment possible.

We make two restrictions on the CCT environment. First, we assume throughout that household has no advanced knowledge of the CCT program before it has begun. If a household with a child of age $t - s$, $s \geq 1$ knew that a CCT program would be implemented when the child was of age t , it would, in general, alter its investments so as to be able to better take advantage of the future CCT. Due to our announcement timing assumption, there are no possibilities of strategic decision-making by the household.

Second, we also limit attention to the case of a program that exists only for one period. The multi-period case brings the possibility of strategic behavior into play, as the analyses of Weitzman (1990) and Macartney (2013) demonstrate. While allowing for strategic long-term play by the household could alter the effect of the program on child development, allowing for this behavior makes the modeling exercise much more complex and less transparent. For this reason we ignore multi-period programs in this paper.

Before getting to the specifics, we want to emphasize the special role that a CCT plays under our modeling assumptions. In order to obtain closed-form decision rules for the large number of endogenous variables, DFW (2014) assume exogenous wage and income processes and the absence of capital markets that would allow the household to save or borrow. The existence of a CCT allows the household to transfer resources between periods t and $t + 1$, albeit in a limited manner. If a household satisfies the conditions to receive a cash reward of ϕ_C in period $t + 1$, then it has increased its nonlabor income in that future period to $I_{t+1} + \phi_C$. If the household has to alter its period t decisions to obtain the reward (which means that k_{t+1} must be larger than it otherwise would have been in the absence of the program), then it must trade off period t consumption of market goods and leisure for period $t + 1$ consumption of these goods. Only with a CCT does the household face such a choice under the DFW structure.

It will be helpful to reformulate the household's problem in order to analyze the response to a CCT. We define the function

$$\begin{aligned}
 J_t(\omega_t, k_{t+1}) &= \max_{i_t} \alpha_1 \ln l_{1,t} + \alpha_2 \ln l_{2,t} + \alpha_3 \ln c_t + \alpha_4 \ln k_t \\
 \text{subject to: } k_{t+1} &= f_t(k_t, \tau_{1t}(a), \tau_{2t}(a), \tau_{1t}(p), \tau_{2t}(p), e_t) \\
 T &= l_{jt} + h_{jt} + \tau_{jt}(a) + \tau_{jt}(p), \quad j = 1, 2 \\
 c_t + e_t &= w_{1t}h_{1t} + w_{2t}h_{2t} + I_t,
 \end{aligned}$$

where i_t is the vector of period t decision variables. J_t defines the maximum amount of utility in period t attainable by a household with characteristics ω_t given that it produces k_{t+1} units of child quality. (Recall our timing convention: parental inputs in period t produce child quality k_{t+1} at the start of period $t + 1$.) When we examine common CCTs,

which set precise levels of k_{t+1} required for the household to earn the reward, we will use J_t to determine the maximal amount of utility the household could obtain in period t if it exactly met the requirement of the CCT. We can return to the original household problem given in (3), since

$$V_t(\omega_t, P_0) = \max_{k_{t+1}} J_t(\omega_t, k_{t+1}) + \beta V_{t+1}(\omega_{t+1}, P_0)$$

and

$$k_{t+1}^* = \arg \max_{k_{t+1}} J_t(\omega_t, k_{t+1}) + \beta V_{t+1}(\omega_{t+1}),$$

where ω_{t+1} includes the child quality level k_{t+1} . k_{t+1}^* is the optimal level of child quality produced under the baseline policy P_0 , i.e. in the “unconstrained” case in which there is no CCT. A further description of the J_t function in our modeling framework is provided in Appendix B.

3.3.1 Level-Based CCT

A level-based CCT policy, P_L , transfers an amount ϕ_L to all households with a child of age $t + 1$ who has attained a development level of k_L . As pointed out in the previous subsection, for a household in which the optimal child quality k_{t+1}^* satisfies $k_{t+1}^* \geq k_L$, there is no change in behavior and the household simply receives the additional income of ϕ_L in period $t + 1$ as a rent. There is only a decision to make in households for which $k_{t+1}^* < k_L$. In this case, the household may increase its production of child cognitive ability by the amount $k_L - k_{t+1}^*$ in order to qualify for the transfer of ϕ_L . This involves a utility loss in period t but a gain in the value of the household’s problem in period $t + 1$ due to the reward received. Since the problem is convex, it will never be optimal to produce more than the minimal level required for the award. The household’s other choice is to produce the original amount k_{t+1}^* and forego the transfer of ϕ_L .¹⁰

Consider a household which, under the baseline policy, would produce less than the required level of child quality: $k_L > k_{t+1}^*$. Using the framework described above, the value to these household of increasing child investments to meet the requirement of the CCT is given by

$$V_t(\omega_t, P_L) = J_t(\omega_t, k_L) + \beta V_{t+1}(\omega'_{t+1}, P_L),$$

where ω'_{t+1} are the updated state variables given the household has meet the CCT program requirement. In this case the household receives the reward ϕ_L and non-labor income is given by $I'_{t+1} = I_{t+1} + \phi_L$. The new level of child development in this case is $k'_{t+1} = k_L$.

¹⁰This is similar to many policies in which some agents, even in the absence of the policy, already comply with the minimum standards set forth in the policy. One example is the case of child support transfers analyzed in Del Boca and Flinn (1995). In that case, a noncustodial father who would voluntarily transfer amount t to the custodial mother faced a child support order of s . If $t < s$, the father would either transfer exactly s and avoid a penalty of θ or would transfer the amount t and pay the noncompliance cost of θ .

The household will respond to the CCT by producing more child quality when

$$\begin{aligned}
V_t(\omega_t, P_L) &\geq V_t(\omega_t; P_0) \\
\Rightarrow J_t(\omega_t, k_L) + \beta V_{t+1}(\omega'_{t+1}, P_L) &\geq J_t(\omega_t, k_{t+1}^*) + \beta V_{t+1}(\omega_{t+1}, P_0) \\
\Rightarrow V_{t+1}(\omega'_{t+1}, P_L) - V_{t+1}(\omega_{t+1}, P_0) &\geq \beta^{-1}(J_t(\omega_t, k_{t+1}^*) - J_t(\omega_t, k_L)).
\end{aligned} \tag{4}$$

We know that the expressions on both sides of the last line are positive for households in which $k_L > k_{t+1}^*$. The left hand side is the size of the gain in the household's problem in the next period. The right hand side measures the utility loss in period t inflated by the inverse of the discount factor, and this expression is bounded from below by 1. It is interesting to note the important role of the discount factor in the decision to respond to the CCT: more patient households are more likely to respond to the CCT program, other things equal.

In summary, for a particular CCT program, there will exist three types of households: (i) households that will earn the reward with no modification of their behavior; (ii) households that will not modify their behavior and will not receive the reward; and (iii) households that will modify their behavior to meet the reward criteria. The set of households receiving the reward consists of those households for which the baseline level of cognitive ability already meets the threshold and those households that change their behavior from the baseline level in order to receive the reward. The relative size of the set of households that respond to the CCT and the distribution of the child quality increases in these households are measures of the success of the CCT. We will consider these outcome measures in Section 4.

3.3.2 Change-Based CCT

This type of CCT, P_G , where the G subscript is meant to connote “growth,” is perhaps the most often used, at least in CCTs that are based on measured outcomes. The CCT consists of a growth-level requirement and a payment. We will denote the fixed payment amount associated with P_G by ϕ_G . The relative change requirement for the award is given by ρ_G , so that any household for which $k_{t+1}^*/k_t \geq \rho_G$ receives a transfer in period $t + 1$ of ϕ_G . Note that given our no anticipation assumption, the current period level of child development k_t is fixed; the household can only alter its investments in the current period and affect next period's level of child development k_{t+1}^* .

As with the level based CCT, for those households that already meet the target, those for which $k_{t+1}^*/k_t \geq \rho_G$, the CCT does not change investment behavior or the level of child quality. These households receive an income transfer of ϕ_G in period $t + 1$ that is essentially an unrestricted transfer. All other households face the decision of foregoing the transfer or increasing the level of child quality to $\rho_G k_t$ to receive ϕ_G in period $t + 1$. The value of receiving the transfer for these households is given by

$$V_i(\omega_t, P_G) = J_t(\omega_t, \rho_G k_t) + \beta V_{t+1}(\omega'_{t+1}, P_G),$$

where $k'_{t+1} = \rho_G k_t$ and $I'_{t+1} = I_t + \phi_G$.

From this expression, we can see the importance of the assumption that the program is only announced when the child is age t (no advanced knowledge). If the program for children of age t was announced when the child was age $t - 1$, for example, the household would have an incentive to decrease the ability level of the child at age t in order to make the value-added criterion easier to satisfy. This is the essence of the ratchet effect explored by Macartney (2013) in his school quality application.

The decision of whether to meet the program requirements is given by (4) as in the level-based case. In other words, with a hard threshold for satisfaction of the CCT requirement, households that do not automatically qualify for the award by their baseline decision only have to consider the value of altering investment decisions enough to exactly meet the criterion for payment.

3.3.3 Piece-Rate CCT

Although many CCT programs have complex reward systems providing increasing rewards for increasingly greater household investments or outcomes, to our knowledge, a pure, or even closely approximate, continuous piece rate CCT has yet to be implemented. By a piece-rate CCT, we mean one for which a price ϕ_{PR} is paid for each unit of k^*_{t+1}/k_t . Then, if $k^*_{t+1} = k_t$, the no growth case, the household would receive a payment of $\phi_{PR} \times 1 = \phi_{PR}$. Positive growth ($k^*_{t+1}/k_t > 1$) results in a payment greater than ϕ_{PR} , and depreciation ($k^*_{t+1}/k_t < 1$) results in a payment less than ϕ_{PR} .^{11, 12} In the pure continuous payoff system we consider here, even households that experience decreases in the level of child ability receive some reward for lowering the level of depreciation, even if they do not achieve positive growth.

Unlike the level-based and change-based CCTs, that set a single threshold for obtaining the reward, the reward in the piece-rate system is a continuous function of the k^*_{t+1} . Thus, the presence of the CCT affects the first order conditions that determine k^*_{t+1} directly, and clearly increases the value of k^*_{t+1} in period $t + 1$ over and above what it was in the absence of P_{PR} . This implies that the production of child ability will be greater in period $t + 1$ for all households. In the threshold-based CCTs examined above, not all households increased investments; the households that “complied” were those that had a planned level of child quality below the threshold without the program but increased their investments in child quality so as to attain the required level to qualify for the reward. The piece rate CCT program is distinguished by two features: (i) the response rate to the program is

¹¹Recall that in our technology with the parameter associated with current child quality δ_{4t} less than 1, low levels of investment can lead to the child regressing and the stock of child skills depreciating.

¹²We use this payoff scheme instead of others that may seem a bit more natural, such as $\max[\phi_{PR} \left(\frac{k^*_{t+1}}{k_t} - 1 \right), 0]$, because we want to examine a system that achieves a positive response from all households. When only improvements in child quality earn rewards, the payoff function has a discontinuity at 0.

100 percent, and (ii) all households receive rents from the program since the households are rewarded for the level of child quality they would have produced in the absence of the program, k_{t+1}^* . Formally, the problem is defined as

$$V_t(\omega_t, P_{PR}) = \max_{k_{t+1}} J_t(\omega_t, k_{t+1}) + \beta V_{t+1}(\omega'_{t+1}, P_{PR}),$$

where the state variables are updated to reflect the piece rate reward added to non-labor income: $I'_{t+1} = I_{t+1} + \phi_{PR}(k_{t+1}/k_t)$.

3.4 Extensions of the Basic Model

In describing the CCTs above, we have ignored a number of practical issues in their implementation and household responses to these programs. Perhaps the most questionable assumption we have made is that the child quality level is perfectly observable by the institutional agent. We explore the impact of relaxing this assumption on the investment decisions of households. Another strong assumption is that the household has total control over the production process, in the sense of knowing, in a deterministic sense, the relationship between its input choices and next period's realized value of child quality. We will also examine how randomness in the production process affects household decision making under a CCT. For simplicity, and due to space limitations, we will only consider the case of a threshold level-based CCT for both of the generalizations that we examine.

3.4.1 Measurement Error Case with a Level Threshold-Based CCT

As above, we assume that an unannounced CCT program is initiated for children of age t . The families of all students whose measured child quality at age $t + 1$, \tilde{k}_{t+1} , exceeds the level k_L will be paid an amount ϕ_L . The measured \tilde{k}_{t+1} is related to the true (latent) child quality k_{t+1} by

$$\tilde{k}_{t+1} = k_{t+1} \times \exp\{\varepsilon_{t+1}\},$$

and for simplicity we assume that ε_{t+1} is distributed as $N(0, \sigma_\varepsilon^2)$, with $0 < \sigma_\varepsilon^2 < \infty$. We can rewrite the condition for a reward as

$$\ln \tilde{k}_{t+1} \geq \ln k_L.$$

Then the probability that the reward is obtained is given by

$$\begin{aligned} P(\ln \tilde{k}_{t+1} \geq \ln k_L) &= P(\ln k_{t+1} + \varepsilon_{t+1} \geq \ln k_L) \\ &= P(\varepsilon_{t+1} > \ln k_L - \ln k_{t+1}) \\ &= \tilde{\Phi}((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon), \end{aligned}$$

where $\tilde{\Phi} \equiv 1 - \Phi$. From this expression, we see that any household has a positive probability of receiving the reward ϕ_L , but the probability of receiving the reward is strictly increasing

in $\ln k_{t+1}$. The value of receiving the reward is the change in next period's nonlabor income, from I_{t+1} to $I_{t+1} + \phi_L$. Then the probability that the household's period $t + 1$ nonlabor income is I_{t+1} is $\Phi((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon)$, while the probability that nonlabor income is $I_{t+1} + \phi_L$ is $\tilde{\Phi}((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon)$.

Returning to household problem in DFW, the key component of the household's decision to invest in the child is the marginal value of child quality to the household in next period. In Appendix A we show that with no CCT program (under the baseline policy P_0) this is given by

$$\frac{\partial E_t V_{t+1}(\omega_{t+1})}{\partial \ln k_{t+1}} = \eta_{t+1}, \quad (5)$$

where η_{t+1} is solved through a backwards recursion from the terminal period M to the current period t , and involves the discount rate, the household's preferences over child quality, and the child development technology. The household's optimal investments in children through time or goods in period t is increasing in η_{t+1} . As shown in Appendix A, η_t is independent of I_t for all t .

Under the CCT program implemented in period t , the state variables ω_{t+1} are altered so that non-labor income is increased to $\omega'_{t+1} = I_{t+1} + \phi_L$ if the household makes sufficient investments in period t so that child quality $k_{t+1} \geq k_L$. Assuming that I_{t+1} is known, along with the other household characteristics in period $t + 1$ except for the realization of the random variable ε_{t+1} ,

$$E_t V_{t+1}(\tilde{\omega}'_{t+1}) = \{\Phi((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon)V_{t+1}(\underline{\omega}'_{t+1}) + \tilde{\Phi}((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon)V_{t+1}(\bar{\omega}'_{t+1})\},$$

where $\underline{\omega}'_{t+1}$ is the low value future state in which $I'_{t+1} = I_t$ and $\bar{\omega}'_{t+1}$ is the high value future state in which $I'_{t+1} = I_{t+1} + \phi_L$; of course, in both cases next period's actual child quality level is equal to k_{t+1} .

Under the CCT program with measurement error, the marginal value of an increase in $\ln(k_{t+1})$ is now

$$\begin{aligned} \frac{\partial E_t V_{t+1}(\omega'_{t+1})}{\partial \ln k_{t+1}} &= -\sigma_\varepsilon^{-1} \phi((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon) V_{t+1}(\underline{\omega}'_{t+1}) \\ &\quad + \Phi((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon) \frac{\partial V_{t+1}(\underline{\omega}'_{t+1})}{\partial \ln k_{t+1}} \\ &\quad + \sigma_\varepsilon^{-1} \phi((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon) V_{t+1}(\bar{\omega}'_{t+1}) \\ &\quad + \tilde{\Phi}((\ln k_L - \ln k_{t+1})/\sigma_\varepsilon) \frac{\partial V_{t+1}(\bar{\omega}'_{t+1})}{\partial \ln k_{t+1}}. \end{aligned}$$

Since as above $\partial V_{t+1}(\tilde{\omega}'_{t+1})/\partial \ln k_{t+1} = \eta_{t+1}$ for any value of I_{t+1} , this expression becomes

$$\begin{aligned} \frac{\partial E_t V_{t+1}(\omega'_{t+1})}{\partial \ln k_{t+1}} &= \sigma_\varepsilon^{-1} \phi((\ln k_L - k_{t+1})/\sigma_\varepsilon) \{V_{t+1}(\bar{\omega}'_{t+1}) - V_{t+1}(\underline{\omega}'_{t+1})\} + \eta_{t+1} \\ &> \eta_{t+1}, \end{aligned}$$

where the inequality is strict since $\sigma_\varepsilon > 0$ and with a CCT offering a positive reward $\{V_{t+1}(\bar{\omega}'_{t+1}) - V_{t+1}(\underline{\omega}'_{t+1})\} > 0$. In this case, the discounted expected value of future child quality is greater for all households than it is in the baseline, that is, in the absence of the program. We conclude that a CCT policy still weakly increases the incentive for investment in children relative to baseline even with measurement error. For an arbitrarily large σ_ε , the impact of k_{t+1} on the likelihood of receiving the reward ϕ_L is arbitrarily small, which means that in the limit (as $\sigma_\varepsilon \rightarrow \infty$) the discounted marginal valuation of $\ln k_{t+1}$ converges to η_{t+1} . As $\sigma_\varepsilon \rightarrow \infty$, the CCT becomes a pure lottery and it has no impact on investment.

While there is a clear incentive effect of a CCT policy relative to baseline even in the presence of finite variance measurement error, there is still the question of how the incentive effects change between a no measurement error threshold-based CCT and the same CCT with some measurement error. In order to understand how measurement error affects the household responses to the threshold-based CCT program, we reconsider the three types of households described above. There was the set of households which qualify for the award without changing their planned investment level, i.e., $k_{t+1}^* \geq k_L$. With measurement error, some of these households now have an incentive to increase their investment as there is some chance that without a higher investment they will not receive the reward. There was another set of households that did not meet the threshold level at baseline ($k_{t+1}^* < k_L$) and for which the take-up condition was not met (4). In the no measurement error case, these households do not alter their planned investment behavior. As was true in the first group of households, some of these households have an incentive to increase investment because there is now some chance that with additional investment they might in fact receive the transfer. The final group of households had $k_{t+1}^* < k_L$ at baseline and the take-up condition (4) was satisfied in the no measurement case. In the no measurement error case, these complier households were the only households to actually change their behavior in response to the CCT program. With measurement error, some of these households may have less incentive to invest as there is now some uncertainty as to whether their investments will be sufficient for child quality to reach the threshold.

The upshot of this discussion is that the average amount of child quality in period $t+1$ could be *increased* if the social planner actually introduced unbiased measurement error into the testing procedure for earning the CCT reward in period $t+1$. Whether or not this is the case, and, if so, what the optimal level of σ_ε is, depends on the actual values of the parameters of the model.

3.4.2 Uncertainty in the Production Process

As before, we consider the simple level threshold case, and we will show that the results obtained are essentially isomorphic to the measurement case just considered given the structure of our model. Recall that the production process is given by

$$k_{t+1} = R_t \hat{k}(i_t, k_t; \delta_t),$$

where R_t is TFP for children of age t and \hat{k}_t contains all of the other components of the deterministic production process. To introduce uncertainty, we assume that R_t is only revealed at the end of period t , after all of the inputs i_t have been chosen. Thus, the inputs i_t are chosen without knowledge of what the final level of child quality will be in period $t + 1$. For simplicity and to facilitate comparison with the measurement error case, let

$$R_t = \bar{R}_t \exp(\varepsilon_t),$$

and assume that $\bar{R}_t = 1$, so that $\ln R_t = \varepsilon_t$, with ε_t distributed as a $N(0, \sigma_\varepsilon^2)$. The household earns the reward when

$$\ln k_{t+1} \geq \ln k_L,$$

or

$$\begin{aligned} \ln \hat{k}_t + \varepsilon_t &\geq \ln k_L \\ \varepsilon_t &\geq \ln k_L - \ln \hat{k}_t. \end{aligned}$$

Unlike the case just analyzed, the uncertainty regarding R_t also impacts next period's welfare through the random amount of child quality produced. The expected value of next period's problem is given by

$$\begin{aligned} E_t V_{t+1}(\tilde{\omega}'_{t+1}) &= \int_{\ln k_L - \ln \hat{k}_t}^{\ln k_L - \ln \hat{k}_t} V_{t+1}(\underline{\omega}'_{t+1}(\varepsilon)) d\Phi\left(\frac{\varepsilon}{\sigma_\varepsilon}\right) \\ &\quad + \int_{\ln k_L - \ln \hat{k}_t}^{\ln k_L - \ln \hat{k}_t} V_{t+1}(\bar{\omega}'_{t+1}(\varepsilon)) d\Phi\left(\frac{\varepsilon}{\sigma_\varepsilon}\right), \end{aligned}$$

where $\underline{\omega}'_{t+1}(\varepsilon)$ includes $I'_{t+1} = I_{t+1}$ and $k_{t+1} = \exp(\varepsilon_t) \hat{k}_t$ and where $\bar{\omega}'_{t+1}(\varepsilon)$ contains $I'_{t+1} = I_{t+1} + \phi_L$ and $k_{t+1} = \exp(\varepsilon_t) \hat{k}_t$. Using Liebnitz's rule,

$$\begin{aligned} \frac{\partial E_t V_{t+1}}{\partial \ln \hat{k}_t} &= -V_{t+1}(\underline{\omega}'_{t+1}(\ln k_L - \ln \hat{k}_t)) \sigma_\varepsilon^{-1} \phi\left(\frac{\ln k_L - \ln \hat{k}_t}{\sigma_\varepsilon}\right) \\ &\quad + \int_{\ln k_L - \ln \hat{k}_t}^{\ln k_L - \ln \hat{k}_t} \eta_{t+1} d\Phi\left(\frac{\varepsilon}{\sigma_\varepsilon}\right) \\ &\quad + V_{t+1}(\bar{\omega}'_{t+1}(\ln k_L - \ln \hat{k}_t)) \sigma_\varepsilon^{-1} \phi\left(\frac{\ln k_L - \ln \hat{k}_t}{\sigma_\varepsilon}\right) \\ &\quad + \int_{\ln k_L - \ln \hat{k}_t}^{\ln k_L - \ln \hat{k}_t} \eta_{t+1} d\Phi\left(\frac{\varepsilon}{\sigma_\varepsilon}\right), \end{aligned}$$

where the results on the second and the fourth line follow from $\frac{\partial V_{t+1}}{\partial \ln \hat{k}_t} = \frac{\partial V_{t+1}}{\partial \ln k_{t+1}} \frac{\partial \ln k_{t+1}}{\partial \ln \hat{k}_t} = \eta_{t+1} \times 1$. This expression simplifies to

$$\begin{aligned} \frac{\partial E_t V_{t+1}}{\partial \ln \hat{k}_t} &= \eta_{t+1} + \sigma_\varepsilon^{-1} \phi\left(\frac{\ln k_L - \ln \hat{k}_t}{\sigma_\varepsilon}\right) \\ &\quad \times \{V_{t+1}(\bar{\omega}'_{t+1}(\ln k_L - \ln \hat{k}_t)) - V_{t+1}(\underline{\omega}'_{t+1}(\ln k_L - \ln \hat{k}_t))\} \\ &> \eta_{t+1}. \end{aligned}$$

This indicates that even with production uncertainty, there is still a greater incentive to invest in children with a CCT program than under the baseline with no program. As in the measurement error case, this implies that there will be increased production of child quality in all households compared to the case of no CCT program. As analyzed in the measurement error case, households that did not respond to the threshold based CCT might now have an incentive to increase investment under the uncertainty case. Those households which were “compliers” under the CCT when there was no measurement error, may produce more or less child quality when $\sigma_\varepsilon > 0$. As in the measurement error case, as next period’s child quality becomes purely random, i.e., as $\sigma_\varepsilon \rightarrow \infty$, there is no impact of random TFP on the level of child quality.

4 Comparing Policies

We have considered three broad classes of policies: unrestricted income transfers, restricted child goods transfers, and conditional cash transfers. It is clear that an unrestricted transfer of x dollars increases household utility no less than does a transfer of x dollars of child goods or a conditional cash transfer of x dollars contingent on the household obtaining some child quality threshold. From the perspective of changing the level of child development, the ranking of the performance of policies is much less clear. We begin this section with a general discussion of some of the important factors that should be considered when selecting a policy. We follow this with a cost-based framework within which we can compare all policies, no matter what their specific form and regardless of the measures used for child outcomes. The following section presents numerical results on the policy comparisons based on estimates from DFW.

4.1 Important Considerations in the Implementation of Policies

4.1.1 Household Consumption

Each class of policies allows the household varying scope to consume the transfer rather than using it to make goods or time investments in child development. In terms of an unrestricted transfer of $\phi_U > 0$, the transfer is simply an increase in household non-labor

income and can be used to purchase increased leisure of the parents and household consumption just as any other dollar of household income could be used. Since child quality is a normal good, some of the increased income will be spent on investment in the child (either to fund higher expenditures on child specific goods or parental time with children), which will increase the level of child quality. The proportion of the transfer consumed rather than invested in children depends on the preferences of the household (preferences for parental leisure and consumption relative to child development), household resources (pre-transfer non-labor income and wage offers), and the productivity of investments. Given that the production technology for child development is dynamic and that the productivity of different types of inputs changes with the child's age, the proportion of the transfer invested in children can also vary depending on the age of the child when the transfer is received.

In the case of a restricted transfer ϕ_R , the impact on household consumption is exactly the same as in the case of an unrestricted transfer when the household was planning to spend at least the amount of ϕ_R on child investment. When a household's planned spending on child goods is significantly less than ϕ_R , the restricted transfer distorts the household's decisions away from the consumption of market goods and leisure and results in increased child investment, relative to the case of an unrestricted transfer. The restricted transfer therefore results in less household consumption of the transfer overall because for some households the restricted transfer is "binding." The proportion of households for which the restricted transfer is binding depends on the current level of investment and the primitives that give rise to this expenditure (preferences, technology, and resources). In general, households that are resource poor or that value parental leisure and consumption much more than child development would have small baseline child goods expenditures and much of the restricted transfer could not be consumed. With the same transfer amount in both cases, the restricted transfers will result in a distribution of child quality that dominates the unrestricted transfer case, and the consumption of leisure and market goods will be (weakly) less than in the unrestricted transfer case.

In the case of CCTs, if the household responds to a CCT offered in period t (with reward received in period $t + 1$), there will be increased investment in the child, and consequently household consumption of leisure and market goods will fall in that period. This is offset by the household's gain in consumption of all goods in the next period when the reward is received. For households that already meet the performance criteria (as in the threshold-based CCTs), no adjustment in investment is induced by the transfer program, and the CCT is simply an unrestricted transfer for these households, although the transfer is received the following period rather than in the current period. Because there are in general some "compliers" with the CCT program, there is less household consumption of the transfer than in the unrestricted case, and one would expect the CCT program to dominate the unrestricted transfer in its effect on child development at a given cost.

4.1.2 Input Distortion

Of the three classes of policies we have considered, the only one that potentially leads to distortions in the utilization of inputs is P_R , the restricted transfers case. This is due to the fact that some households, with relatively low expenditures on “child goods” in the absence of the program, are given a large amount of such goods under the transfer policy. This causes them to produce child quality in an inefficient manner. Of course, the problem is that when given the unrestricted transfer, the household would produce less child quality than it would when giving the household a restricted transfer. This is the tradeoff the institutional agent faces in terms of P_U and P_R .¹³ The CCT case, P_C , is beneficial in that the household is free to determine its input mix, but must self-finance the investment in period t so as to receive the reward in period $t + 1$.¹⁴

4.1.3 Measurement Uncertainty

Except for a brief consideration of the impact of mis-measurement of child development in a simple threshold-based CCT case, we have not discussed measurement issues. While measurement of the child’s level of development before and after the implementation of any program is crucial for investigating its effectiveness, in the case of P_U and P_R , the researcher may be content with the use of “noisy,” yet unbiased, measures at the individual level if only estimates of aggregate, population wide, performance is required. However, the measurement of performance is crucial to the incentive effects of a CCT.

As we observed in the DFW model, even with classical measurement error,¹⁵ a CCT still increased the amount of child quality produced. While this is a positive result from the standpoint of an institutional agent interested in increasing child quality, it also obviously

¹³In general, unrestricted transfers do not distort the input mix in the sense that the household can use the income transfer to optimally allocate resources toward child goods expenditures or to “purchase” more time with children by reducing labor supply. However, an important feature of our model is that we allow households to choose a corner solution in labor supply, where one or more parents is not working at all. This aspect of our model matches the empirically relevant pattern that many mothers with young children are out of the labor force. In the case in which a parent is supplying 0 labor hours, the household cannot use unrestricted income transfers to reduce this parent’s labor supply and free up time for investment in children. This implies that the unrestricted transfer distorts the input allocation toward goods and away from time in the same way as restricted transfer, although the distortion is considerably less in the unrestricted case because it is only distortionary for the minority of households at the 0 labor supply corner at baseline.

¹⁴It should be noted that the household is not necessarily choosing the child development maximizing mix of inputs. As DFW (2014) demonstrate, the household is choosing the household optimal mix of inputs taking into consideration the particular leisure and consumption costs for that household. For example, households with a high preference weight on father’s leisure can be assigning too little of the father’s time to child development relative to the technologically-determined productivity of the father’s time.

¹⁵By classical measurement error, we mean the case in which the measurement error is independently and identically distributed across individuals and over time with mean 0 and constant variance σ_ϵ^2 . We assumed this applied to $\ln k$ measurement, and added the assumption of normality.

involves some welfare loss on the part of households. Moreover, as the measurement error becomes increasingly large, the impact of the CCT on child investment decisions lessens.

In the case of a change-based CCT, measurement issues become even more serious due to the fact that readings on child ability are required at two points in time. Imagine that the household knows the child's true ability at both times t and $t + 1$, which is a requirement of our model, but that the institutional agent only observes $\ln \tilde{k}_t$ and $\ln \tilde{k}_{t+1}$, and gives the reward ϕ_G when $\ln \tilde{k}_{t+1} - \ln \tilde{k}_t \geq \ln \rho_G$. Consider the case of two otherwise identical households, where in one $\ln \tilde{k}_t \ll \ln k_t$ and in the other $\ln \tilde{k}_t \gg \ln k_t$, meaning that in the first household the measurement error draw was negative and large in absolute value while in the second the measurement error draw was large and positive. In the first household, investment in child quality will be lower than in the second household, since the likelihood of satisfying the growth requirement is easier when the observed baseline value is lower. Clearly, if households could manipulate the measurement error process, such as by instructing the child to under-perform on the period t test, it would be in their interest to do so. These types of considerations are extremely important, since output-based CCTs will clearly rely on measurements containing random errors that lead to strategic behavior on the part of the household.

4.1.4 Administrative Costs

In theoretical investigations of optimal policies, such as tax systems in public finance (e.g., Mirrles, 1971), no consideration is given to the costs of implementation of these policies. In terms of the transfer systems we are considering here, there are clearly very different levels of administrative costs associated with each.

The least costly transfer system from an administrative standpoint is P_U . To make these transfers it is only necessary to determine the eligibility of each household in the population, which in our simple case just involves verifying that there is a child of age t living there. As we will see below, the P_U policy is easily dominated by the P_R and P_C ones in terms of standard performance measures, but when one adds administrative costs into the competition, their dominance is likely to be significantly reduced.

The costs of the restricted transfers program, P_R , would mainly consist of providing transfers that could only be used in child production. If this would involve establishing a child development intervention program (such as Perry pre-school or Abcederian-type program), there are clearly large costs associated with providing such direct interventions in terms of teachers (or case workers) and buildings. Even a program involving less intervention, such as the provision of tuition subsidies or other child goods directly to the household, may involve considerable outlays in ensuring that these goods were only used by the intended households in the appropriate manner.

The CCT programs transfer only money directly to the household, so there is no need to monitor how these transfers are spent. There are likely to be significant costs in designing an effective CCT and in assessing whether performance standards have been met.

4.2 A Metric for Comparing Policy Performance

Although we have discussed some factors that are important in evaluating the performance of programs, it remains to provide a relatively comprehensive measure of policy performance that can be used to compare their efficacy. We provide a formal way to compare these policies in this section.

The first order of business is the specification of the objective function of the institutional agent. For purposes of simplicity, we will assume that the planner's objective is defined in terms of average child quality in the population of interest. It is straightforward to utilize other criteria, such as the minimization of the interquartile range, or even an entire set of criteria. In terms of the group of interest to the planner, it may be the entire population of households with a child of age t , or a targeted subset of the population defined in terms of observable characteristics, such as the income level of the household or whether it is headed by a single female.

Denote the current policy environment (current taxes and transfer policy) facing all households as P_0 . Consider the announcement of policy P_j for children of age t to be implemented immediately and to last only during the year in which the child is age t . If the policy would not have been implemented, the average level of child quality produced while the child was age t would have been

$$\int k_{t+1}(\omega_t; P_0) dF_{\Omega_t}(\omega_t).$$

where $F_{\Omega_t}(\omega_t)$ is the population distribution of household characteristics. Under the transfer policy P_j , the average child quality produced is instead

$$\int k_{t+1}(\omega_t; P_j) dF_{\Omega_t}(\omega_t).$$

We will take the planner's objective as being the improvement of average quality, which under policy P_j is

$$\frac{\int k_{t+1}(\omega_t, P_j) dF_{\Omega_t}(\omega_t)}{\int k_{t+1}(\omega_t, P_0) dF_{\Omega_t}(\omega_t)}. \quad (6)$$

Our benchmark evaluation criteria is 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. We can think of the planner's problem as consisting of two stages. In the first, for a given improvement level of \hat{Z} and a policy "type" of P_j , the planner finds the best set of policy parameters to achieve that pre-specified gain. Let the parameters characterizing policy P_j being summarized by the vector θ_j . Then the first step of the planner's problem is

$$\begin{aligned} & \min_{\theta_j} C_j(\theta_j) \\ \text{subject to } \hat{Z} &= \frac{\int k_{t+1}(\omega_t, P_j) dF_{\Omega_t}(\omega_t)}{\int k_{t+1}(\omega_t, P_0) dF_{\Omega_t}(\omega_t)}, \end{aligned}$$

where $C_j(\theta_j)$ are the total monetary costs (solely in terms of transfers made) associated with the transfer policy P_j . The solution to this problem is denoted $\theta_j^*(\hat{Z})$, and the associated costs are $C_j^*(\hat{Z}) = C_j(\theta_j^*(\hat{Z}))$. The final stage of the problem is to choose the best policy from the set of options available, where we denote the set of feasible options by φ . Then the minimum cost way to obtain a gain of \hat{Z} in the planner's objective is given by

$$C^{**}(\hat{Z}) = \min_{P_j \in \varphi} \{C_j^*(\hat{Z})\},$$

and the optimal policy that attains this minimum cost is denoted $P^{**}(\hat{Z})$.

Now we consider some examples of the use of this framework. We begin with the simplest case, that of an unrestricted transfer, P_U . Say that the planner desires a two percent gain in the average ability in the population compared with baseline, so that $\hat{Z} = 1.02$. The policy P_U is characterized only by the transfer level of ϕ_U , which is given to all households with a child of age t , so that $\theta_U = \phi_U$. Let there be N individuals in the population, so that $C_U(\theta_U) = N \times \phi_U$. Since child quality is a normal good in the household, an increase in the household's nonlabor income by the amount ϕ_U leads to an increase in child quality for each household, so that the planner's objective (6) is continuously increasing in ϕ_U . This means that there is a transfer amount $\phi_U^*(\hat{Z})$ that exactly attains the target of \hat{Z} , with any $\phi_U > \phi_U^*(\hat{Z})$ producing a larger gain in average quality but at a higher price. Then the minimum cost of the unrestricted transfer program to obtain the level \hat{Z} is $C_U^*(\hat{Z}) = N \times \phi_U^*(\hat{Z})$.

The computation of the cost of the restricted transfer program is similarly straightforward. In this case, the cost is given by $C_R(\theta_R) = N \times \phi_R$, since all households receive a child goods bundle with a value of ϕ_R dollars. As was true under P_U , average child quality levels are strictly increasing in the transfer ϕ_R , both due to the income effect and because the transfer of child goods causes a distortion in household consumption decisions that results in no less child quality being produced than under an unrestricted transfer of the same amount. This implies that $\phi_R^*(\hat{Z}) \leq \phi_U^*(\hat{Z})$ for all \hat{Z} , and implies

$$C_R^*(\hat{Z}) \leq C_U^*(\hat{Z}), \text{ for all } \hat{Z}.$$

Taken at face value, this result implies that the unconditional transfer policy is weakly dominated by a restricted transfer policy, so that the institutional agent should never implement P_U if their objective only involves improvements in average child quality. If these are the only two policy choices available to the planner and if P_U is chosen, this may indicate that the objective of the planner is not what we have taken it to be or that the costs of implementation of P_R are sufficiently great that the difference $C_U^*(\hat{Z}) - C_R^*(\hat{Z})$ does not offset the differences in the cost of implementation.

In evaluating the efficiency of a CCT, matters are slightly more subtle. For simplicity, we will only consider threshold-based CCTs, those defined in terms of absolute levels or growth rates. As was discussed above, the payoff to the CCT depends not only the

characteristics of the household as summarized by ω_t , but also the state variables of the period in which the payoff occurs, which are elements of the vector ω_{t+1} . This is because the value of the payoff when the child is age $t + 1$ depends on the wage rates of the parents and the nonlabor income of the household in $t + 1$. Then we have to redefine the objective of the institutional agent to be

$$\frac{\int k_{t+1}(\omega_t, \omega_{t+1}, P_C) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})}{\int k_{t+1}(\omega_t, P_0) dF_{\Omega_t}(\omega_t)}, \quad (7)$$

since that the joint distribution of the state variables in period t and $t + 1$ is required to assess the child quality improvement under the CCT.

For a given household characterized by the characteristic vectors (ω_t, ω_{t+1}) , we will be able to determine whether it responds to P_C using the condition given in (4). In the case of a level-based award, the set of households that receive the reward consists of those for which $k_{t+1}^* \geq k_L$, which means that they receive a pure “rent” from the existence of the program and do not alter their period t behavior, and the subset of households for which $k_{t+1}^* < k_L$ but for which the take-up condition is satisfied. Let $d(\omega_t, \omega_{t+1}; k_L, \phi_L) = 1$ if the household characterized by (ω_t, ω_{t+1}) receives the transfer ϕ_L when the child is aged $t + 1$ given the threshold requirement of k_L , and set $d = 0$ if this is not the case. Similarly, let $d_1(\omega_t, \omega_{t+1}; k_L, \phi_L) = 1$ if household (ω_t, ω_{t+1}) “automatically” qualifies for the transfer ϕ_L because $k_{t+1}^* \geq k_L$, and let $d_2(\omega_t, \omega_{t+1}; k_L, \phi_L) = 1$ if household (ω_t, ω_{t+1}) has $k_{t+1}^* < k_L$ but increases k_{t+1}^* to k_L to receive the transfer ϕ_L when the child is aged $t + 1$. In this case, the entire cost of the program is

$$\begin{aligned} C_L(k_L, \phi_L) &= \phi_L N \int d(\omega_t, \omega_{t+1}; k_L, \phi_L) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1}) \\ &= \phi_L N \left\{ \int d_1(\omega_t, \omega_{t+1}; k_L, \phi_L) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1}) \right. \\ &\quad \left. + \int d_2(\omega_t, \omega_{t+1}; k_L, \phi_L) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1}) \right\}. \end{aligned}$$

As before, the constrained cost minimization problem involves selecting a level of improvement $\hat{Z} > 1$ and equating (7) to that value. This establishes a set of values of $(k_L, \phi_L)(\hat{Z})$ that are consistent with producing an improvement of exactly \hat{Z} . We then select the optimal combination $(k_L^*, \phi_L^*)(\hat{Z})$ that is the solution to $\min_{k_L, \phi_L \in (k_L, \phi_L)(\hat{Z})} C_L(k_L, \phi_L)$. The cost of the level-based CCT is then given by $C_L^*(\hat{Z}) = C_L(k_L^*(\hat{Z}), \phi_L^*(\hat{Z}))$. The computation of the minimum cost for a change-based CCT proceeds in an exactly analogous manner.

Although the total cost of the policy is the only relevant characteristic in this formulation, it is of interest to know how effective the policy is in actually changing behavior. There are a variety of measures of effectiveness that seem reasonable, and we will briefly explore a few of them. For simplicity, we will continue our discussion with the level-based CCT framework, but all of these measures extend in a straightforward way to the change-based and piece-rate CCT cases.

One simple measure of effectiveness is the proportion of population members that actually changed their investment decisions in order to obtain the reward, which is simply

$$P(k_{t+1} \neq k_{t+1}^* | k_{t+1} \geq k_L) = \frac{\int d_2(\omega_t, \omega_{t+1}; k_L^*(\hat{Z}), \phi_L^*(\hat{Z})) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})}{\int d(\omega_t, \omega_{t+1}; k_L^*(\hat{Z}), \phi_L^*(\hat{Z})) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})}.$$

Since $k_{t+1} = k_L$ for any household that alters its behavior to obtain the transfer in period $t + 1$, an equivalent representation for this probability is $P(k_{t+1} = k_L | k_{t+1} \geq k_L)$. This measure takes values in $[0, 1]$, where 0 indicates that all households received a pure income transfer in period $t + 1$ without modifying their period t behavior and 1 indicates that all individuals receiving the transfer modified their period t behavior.

This measure does not contain information on the extent of behavioral change due to the CCT. To do so, we should adjust for the difference between the gain in the child's ability from the award. We can write the average gain in ability in households receiving the transfer as

$$\begin{aligned} & \{P(k_{t+1}^* > k_L | k_{t+1} \geq k_L) \times 0\} + \{P(k_{t+1} = k_L | k_{t+1} \geq k_L) \\ & \times \frac{\int (k_L - k_{t+1}^*(\omega_t, \omega_{t+1})) d_2(\omega_t, \omega_{t+1}; k_L^*(\hat{Z}), \phi_L^*(\hat{Z})) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})}{\int d_2(\omega_t, \omega_{t+1}; k_L^*(\hat{Z}), \phi_L^*(\hat{Z})) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})}\}. \end{aligned}$$

This is a useful way to express the average gain, since it illustrates the nature of the institutional agents constraints in setting effective policies when a threshold-based CCT is used. The gains in child ability under the program come from the last term in the expression above. As the threshold for obtaining the payment, k_L , is raised, the larger is the gain in child quality. At the same time, for a fixed payment amount, ϕ_L , the smaller is the size of the set of households willing to make the required investment to obtain the reward.¹⁶ The optimal policy in terms of both variables is found when the best mixture of the level required to obtain the reward and the size of the reward is found that results in a gain of \hat{Z} .

We conclude by briefly mentioning the piece-rate case. Since all households produce positive child quality in equilibrium, the piece-rate CCT produces a gain in child quality for all households. In terms of the sets defined above, all households take part in the program, that is, $d(\omega_t, \omega_{t+1}; \phi_{PR}) = 1$ for all (ω_t, ω_{t+1}) . Moreover, all households change the level of their investment with respect to baseline, so that $d_2(\omega_t, \omega_{t+1}; \phi_{PR}) = 1$ for all (ω_t, ω_{t+1}) . The cost of the piece rate system is

$$C_L(\phi_{PR}) = \phi_{PR} \int \frac{k_{t+1}^*(\omega_t, \omega_{t+1}; \phi_{PR})}{k_t} dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1}),$$

¹⁶This is similar to the situation faced by the institutional agent in the child support award analysis of Del Boca and Flinn (1995). In their case, the institutional agent could increase the size of the child support order but faced a decreasing probability that the noncustodial father would comply with the order. The situation in that model was simpler than the case considered here, since the distribution of punishments for failure to comply was taken as exogenous. In the case considered here, the institutional agent has two policy variables available, the level of attainment required to obtain the award and the size of the award.

where we recall that k_t is an element of ω_t , which is why we do not include it explicitly in the numerator in the integrand. The gain in average child ability under the piece-rate CCT is given by

$$\frac{\int k_{t+1}^*(\omega_t, \omega_{t+1}; \phi_{PR}) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})}{\int k_{t+1}^*(\omega_t, \omega_{t+1}; \phi_{PR} = 0) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})},$$

where $k_{t+1}^*(\omega_t, \omega_{t+1}; \phi_{PR} = 0) = k_{t+1}^*(\omega_t)$, which is the baseline case. Because this measure is continuously increasing in ϕ_{PR} , there is a unique ϕ_{PR} for every \hat{Z} such that

$$\hat{Z} = \frac{\int k_{t+1}^*(\omega_t, \omega_{t+1}; \phi_{PR}^*(\hat{Z})) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})}{\int k_{t+1}^*(\omega_t, \omega_{t+1}; \phi_{PR} = 0) dF_{\Omega_t, \Omega_{t+1}}(\omega_t, \omega_{t+1})},$$

with $\phi_{PR}^*(\hat{Z})$ strictly increasing in \hat{Z} . The piece-rate system obtains a response from all household members for any $\phi_{PR} > 0$, and because of this may dominate the other systems that we quantitatively evaluate in the next section.

4.3 Limitations of the Analysis

Before discussing the quantitative results, it will be useful to recognize the limits of the analysis we are able to undertake given the assumptions required to estimate the model. The relaxation of some of the more restrictive features of the model is the focus of our current research in child development.

4.3.1 Short-Run Programs

Throughout our discussion and the policy simulations, we only consider the case of short-run program interventions. We consider the case in which 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. At several points we have emphasized the fact that the existence of multi-period programs will generally result in households behaving strategically in a dynamic sense. This type of behavior is more difficult to analyze, although it is certainly possible to do so within our modeling framework.

Even in the context of CCTs in which behavior in one period could result in additional non-labor income in the subsequent period, we saw that households will consider next period's wages and nonlabor incomes (or at least the distribution of them) in deciding on age t actions. Under our modeling assumptions of no saving and borrowing, the existence of a CCT gave the household one avenue to transfer resources across periods. In this sense, the household behaves in a strategic manner, although if the model allowed (at least) savings, the impact of the CCT would be different than what we determine it to be under our restrictive assumptions on capital markets.

4.3.2 Multiple Development Outcomes

Another limitation of our analysis is that our model includes only a single child development outcome k_t . In general, there might be several dimensions of child quality at any given child age, k_{1t}, k_{2t}, \dots , with each outcome produced by a potentially distinct child development production technology, distinct household preferences over each quality dimension, and scope for parents to make assignable investments in each type of “ability.” For example, in the technology estimated by Cunha et al. (2012) there are two child quality levels, cognitive and non-cognitive skills, but the investments are assumed to be the same in both. In terms of evaluating transfer policies, any given policy might have differential effects on the various outcomes given differences in technologies and preferences. In the restricted goods transfer case, assignable goods which are productive in producing cognitive skills but less so in producing non-cognitive skills could have different effects than a policy which provides goods which are equally productive in producing both outcomes. In the CCT case, multiple outcomes substantially complicates the analysis as one now needs to consider how to structure optimally the rewards across multiple outcomes. Furthermore, considering measurement error varying across multiple outcomes, as in the case where cognitive measures are less noisy than non-cognitive measures, could further complicate the analysis as one would need to consider how the measurement error in each measure affects the incentives to invest in children under different types of CCT programs.

4.3.3 Other Actors in the Child Development Process

Our model only includes the investment decisions of parents, with the most important omissions being the child itself, other siblings, and educators. For very young children, parents most probably are the primary investors, but this situation changes dramatically from pre-school on, when teachers and classmates may have a large impact on cognitive and non-cognitive development. From The CDS data, we know that children begin to spend significant amounts of time in self-investment during adolescence (as measured by time spent alone doing homework). Our current research extends the DFW framework to include child-self investment in the production technology, with factor productivity increasing in the age of the child. We model the investment decision of the household as a Stackelberg game pitting the parents (as leaders) against the child (as follower) in making time allocation decisions.

Within our modeling framework, if schooling simply (endogenously) shifts factor productivity in any period (which was denoted R_t), then there is no impact on the child investment decisions of the household, although varying levels of R_t across households in the population could have important effects on the outcome of CCTs. It seems likely that school quality is endogenously chosen by households as a direct input into the production of child quality, and this aspect of household decision-making is captured only by our child goods aggregate but not considered separately.

Both of these generalizations suggest that making transfers to the parents is not the only way in which a policy-maker can positively impact child development. In fact, the results in Behrman et al. (2014). suggest that CCTs that combine rewards to several different agents may be the most effective. In the current framework we only allow transfers to parents, but it is reasonable to expect that CCTs that target students and/or teachers during adolescence may be the most effective way to affect child development.

4.3.4 Targeting Sub-Populations

In many cases, it may be desirable to “target” the transfers, which means that only a subset of the population that is 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, program might be targeted to households with income below some amount. All of the formalism developed above can be easily adopted to the case of targeted transfers or CCTs. Through targeting the planner may be able to greatly reduce the proportion of payments that are ineffective in increasing child quality, at the cost of obtaining improvements in only the targeted sub-populations.

5 Quantitative Results

In this section we examine the simulation results for the three different transfer program types using the point estimates of the DFW model. For the unrestricted and restricted transfer case, we examine a number of characteristics of the program by the size of the transfer and the child’s age at which the transfer is provided. In the final part of this section, we compare the cost effectiveness of the various transfer policies in the manner discussed in Section 4.2. 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. While our theoretical policy analysis is general in the sense that it applies to any one dimensional child outcome, and is not particular to cognitive or non-cognitive skills, our specific quantitative results are based on our measure of child development, the Letter Word score, a measure of cognitive skills, which is discussed in more detail above.

In the quantitative results we compare policies in terms of latent child quality rather than any particular test score measure of child skills, for example. With respect to the rank ordering of policies according to changes in the distribution of child ability, this choice of metric is without loss of generality as any reasonable measure of child development is a monotonic function of the latent child quality level. Comparing the effects of these child development transfer policies externally to other policies requires “anchoring” of the child development level with respect to some interpretable outcome such as labor market earnings or reduction in crime (see Cunha et al. (2012) for a discussion).

5.1 Unrestricted Transfers

Table 3 presents results for three levels of unrestricted transfers: \$100, \$250, and \$500. The time unit is weekly, so that 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. We discuss the distribution of child quality changes within the population later.

5.1.1 Average Effects

Comparing the baseline levels to those with a \$100 transfer we see that the average gain in child quality (6) 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 4.25, about a 2.7 percent gain from baseline. Finally, a transfer of \$500 produces a gain of $4.35/4.15 = 1.04$, or a 4.7 percent gain.

5.1.2 Household Behavior

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 described above) and the productivity of child goods expenditures in producing child quality. For example, even if a household highly values child quality, a monetary transfer to the household at an age for which inputs other than time have limited value in producing child quality will result in the household consuming much of the transfer by increasing its consumption of market goods and leisure.

As shown in Table 3 the unrestricted transfer also reduces labor supply through a standard income effect. The time taken from labor market activity is spent on time investments with the child and on parental leisure. Thus there is a positive spillover effect of income transfers on households: not only do income transfers finance child goods expenditures but also time investments in children. Both parents spend more time in child investment, in both active and passive activities. However, as we see in the increase in parental leisure time, parents also consume a large portion of the transfer through increased leisure. 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, both of which varies by the age of the child.

5.1.3 Timing

Next we explore how the timing of the transfer affects its efficacy. Figure 3 displays the average gain (6) for various levels of unrestricted transfers provided at various child ages. As in all of this analysis, the level of child 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 vertical 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 as the child ages, there are two factors which reduce the effect of unrestricted transfers as the child ages. First, as the productivity of previous period’s child quality stock increases with age, there is less “room for improvement” as the child approaches the end of the development period. This stasis implies any type of intervention will have more limited impacts at the later stages of the development process. Second, the unrestricted transfers not only increase goods allocation to children but also parental time through a reduction in parental labor supply. As our estimates indicate that time with parents is key input into child development and the productivity of this time is declining with age, the positive “spillover” effect is also declining with the child’s age and reducing the effect of the transfer.

5.1.4 Distribution of Effects

In contrast to Table 3, which displays the average level, Figure 4 shows the distribution of effects from various levels of unrestricted transfers given at age 10. 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 (2.5 percent gain), and the average gain at the 90th percentile is 1.083 (8.3 percent gain). The difference in household responses is due to the sources of heterogeneity in the model: wage offers, baseline non-labor income, household preferences, and initial child quality levels. Households with low levels of non-labor income and those with low wage offers but relatively high preferences for child quality respond the most to the transfer. Other households that already have high levels of household non-labor income or relatively low preferences for child quality react relatively little to the program.

5.2 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 is unrestricted. For 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.¹⁷

5.2.1 Average Effects

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 $4.23/4.15 = 1.019$ or 1.9 percent gain compared to a 1 percent gain for the unrestricted transfer. A weekly restricted transfer of \$500 results in over a $4.61/4.15 = 1.111$, or 11.1 percent gain in average child quality from baseline compared to a 4.7 percent gain for an equal transfer in the unrestricted case.

5.2.2 Household Behavior

There are decidedly different household responses to restricted versus 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 increases much more than what we saw in the unrestricted transfer case in Table 3. As the restricted transfer amount increases, the requirement that the transfer be spent on child goods is increasingly binding for more households, and average child good expenditures are increasing rapidly and household consumption is increasing only slightly (as compared to the unrestricted case).

In addition to the goods allocation, there is a smaller effect of the restricted transfer on time investments of parents. Like unrestricted transfers, restricted transfers also have a labor supply effect, even on households where the restriction is binding, as the increase in resources to the household reduces the incentive to work at the margin. However, the labor supply response is less in the restricted case than in the unrestricted case 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 less in the restricted case than in the unrestricted. As a consequence, the increase in time allocated to children is also less in the restricted transfer case.¹⁸

¹⁷The restriction is actually somewhat weaker. The restricted transfer is the same as an unrestricted transfer in households for which an unrestricted transfer of φ_R would lead them to choose to spend more than φ_R on the child.

¹⁸There is an important caveat to this conclusion. In our specification, there is only a single child good,

5.2.3 Timing

Figure 5 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.

5.2.4 Distribution of Effects

Figure 6 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 4), also notable is that there is a much larger right tail reflecting a considerable mass of households that have quite large gains in average child quality as result of the transfer. These highly affected households are those that either have few resources to spend on children but with high preferences for children quality or are resource rich with low preferences for child quality. In the latter case, the in-kind transfer is effectively forcing substantially higher investment in children.

5.3 Comparison of Unrestricted and Restricted Transfers

Figure 7 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 advantage of the restricted transfer is larger as the household consumes a relatively higher proportion of the unrestricted transfer.

5.4 Conditional Transfers

Next we turn to results using conditional transfers. In each of these cases, we consider a “value-added” or growth type criterion 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, ρ_G , and the level of reward incentive, ϕ_G . Any household for which $k_{t+1}/k_t \geq \rho_G$ receives a transfer in period $t + 1$ of ϕ_G .

which includes all goods provided to the child including tutoring, books and toys, and any child care services. If the restricted transfer were to take the form of a subsidy for child care services, which directly substitute for parental time with the child, then it may be possible that for some households this policy would reduce parental time with children and increase labor supply. These type of effects require a richer model with multiple child goods and an explicit accounting for the time budget of the child including time inputs from several agents.

5.4.1 Take-up

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 ρ_G 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_G . 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 $\phi_G = 0$) 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 $\phi_G = 0$), we see the proportion of the population that already satisfies the target at baseline, which is declining in the threshold. Moving across the columns provides the incentive effect of increasing the one-time reward payment in period $t + 1$ (age 11 in this simulation) and fixing the target threshold. As we move along the columns, the take-up rate increases as the reward increases. The fraction of compliers in the population can be calculated by the difference in the take-up rate for some given reward relative to the baseline of 0 reward.

5.4.2 Per Capita Costs

In the unrestricted and restricted cases, the per capita cost of those programs is trivial to calculate since all households receive the given transfer level. In the CCT case, only households that successfully meet the program requirements receive the reward and contribute to the cost. Table 6 provides the per capita cost of the various combinations of conditional transfer policy parameters. The per capita cost is simply the proportion of households that would meet the threshold (from table 5) multiplied by the reward level ϕ_G . In the top row, with a target threshold of 0, the per capita cost is equal to the reward ϕ_G since all households receive the reward. As the threshold target increases as we move down the columns, the per capita cost declines as the proportion of households satisfying the performance criterion declines.

5.4.3 Average Effects

Figure 8 displays the average gain in child quality by the individual household child quality threshold (ρ_G) for four different levels of reward $\phi_G = 20, 50, 100, 200$. Average gain is increasing in both the reward level and the threshold target. At low performance targets (ρ_G small), the difference across reward levels in overall average gain is 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 relative few complier households, as most households would already meet the threshold in the baseline, and higher rewards are a pure rent for most households. However, as the performance target increases, more households are no longer automatically meeting the threshold at baseline, and there is greater scope for rewards to incentivize child investments. At these higher

targets, the proportion of compliers increases with the reward level and more households increase their investment to meet the target. The larger share of compliers combined with a higher target increases the average gain in the children’s human capital for the total population.

5.5 Comparison of Unrestricted, Restricted, and Conditional Transfers

Figure 9 compares all three types of transfers (unrestricted, restricted, and conditional). For the conditional transfers, we calculate the minimum cost combination of performance target and reward to achieve a given average gain (from ??). 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 policy (in terms of cost minimization), 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 to make the decision of which inputs to use, as in the case of the conditional transfer, this distortion is eliminated.

6 Conclusion

We have used the model developed and estimated in DFW to examine in some detail the impacts of various types of transfer policies on child development. The model itself is quite stylized, but does incorporate heterogeneity in household preferences, a dynamic child development production technology, and a large set of inputs into the production process, which the household chooses given a dynamically evolving resource constraint. Within this model we considered three broad classes of transfer policies: unrestricted transfer of income, restricted transfer of child goods, and a conditional (cash) transfer, which was given only to households that met a performance criteria.

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 that would not qualify under baseline must efficiently (in the sense of input utilization) adjust their behavior to satisfy the performance criterion and earn the reward. This reward may have even stronger impacts in the long-term because the

reward (received in the following period) is essentially an unrestricted transfer, which we know that households, to varying degrees, will spend at least partly on their own children.

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 to all households. The main costs of the conditional transfer program are in terms of the measurement of the child's human capital and validation of the household having attained the performance threshold. These costs are not trivial, and involve difficult measurement issues. While we did not provide any quantitative results, we analyzed the consequences of these issues, both in the case where child quality is measured with error by the program administrator and when the household itself is uncertain about the productivity of investments. 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. Depending on the household's attitude toward risk and the level of uncertainty, this may have substantial impacts on the performance of the program.

The main limitation of our results is the restricted environment in which the analysis takes place. The model at present only considers one-child households¹⁹ 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 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. (2014) 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 with the results of these types of field experiments, more cost-effective policies to promote child development can be found.

¹⁹In 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.

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. (2014). “Aligning Learning Incentives of Students and Teachers: Results from a Social Experiment in Mexican High Schools.” Forthcoming, *Journal of Political Economy*.
- [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] Caucutt, E. and Lochner, L. (2012). “Early and Late Human Capital Investments, Borrowing Constraints, and the Family.” NBER *Working Paper* 18493.
- [9] 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.
- [10] Cunha, F. and J. J. Heckman, and S. Schennan.(2010). “Estimating the Technology of Cognitive and Non-Cognitive Skill Formation.” *Econometrica* 78: 883-931
- [11] Cunha, F. (2013). “Investments in Children When Markets Are Incomplete.” *Working Paper*
- [12] Dahl, G. B. and Lochner L. (2012). ”The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit.” *American Economic Review*, 102(5): 1927-56.

- [13] Del Boca D, Flinn C. and M. Wiswall (2014). “Household Choices and Child Development ” *Review of Economic Studies* 81 (January 2014): 137-185.
- [14] Dubois, P., de Janvry, A., and Sadoulet, E. (2012). “Effects on School Enrollment and Performance of a Conditional Transfers Program in Mexico.” *Journal of Labor Economics* 30 (3): 555-89.
- [15] Duncan, G. J., P. A. Morris and C. Rodrigues (2011). Does money really matter? Estimating impacts of family income on young children as achievement with data from random-assignment experiments. *Development. Psychology.* 47(5), 1263-1279.
- [16] 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
- [17] 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.
- [18] 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.
- [19] Heckman, J.J. and Mosso S. (2014), “The Economics of Human Development and Social Mobility” NBER Working Paper 19925
- [20] Heckman J.J. and Kautz (2012) “Hard Evidence on Soft Skills“ NBER Working Paper Working 18121
- [21] 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.
- [22] Loken, K. V., M. Mogstad and M. Wiswall (2012): What Linear Estimators Miss: The Effect of Family Income on Child Outcomes. *American Economic Journal: Applied Economics*, 4, 1-35.
- [23] 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.
- [24] Mirrles, J. (1971) “An Exploration in the Theory of Optimum Income Taxation.” *Review of Economic Studies* 38: 175-208.

- [25] Morris P., Aber L., Wolf S. Berg J. (2012) "Using Incentives to Change How Teenagers Spend Their Time, The Effects of New York City's Conditional Cash Transfer Program" MDRC 2012
- [26] 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, .
- [27] Rawlings, L. B. and Rubio, G. M. (2005). "Evaluating the Impact of Conditional Cash Transfer Programs." *The World Bank Research Observer*, 20: 29-55.
- [28] 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
- [29] Skoufias, 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
- [30] 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.

A Decision Rules for the DFW Framework

We can write the conditional factor demands for child inputs, where we are conditioning on labor supply choices and nonlabor income, as

$$\tau_{1,t}^*(a) = (T - h_{1t}) \frac{\varphi_{1,t}(a)}{\alpha_1 + \varphi_{1,t}(a) + \varphi_{1,t}(p)} \quad (1)$$

$$\tau_{2,t}^*(a) = (T - h_{2t}) \frac{\varphi_{2,t}(a)}{\alpha_2 + \varphi_{2,t}(a) + \varphi_{2,t}(p)} \quad (2)$$

$$\tau_{1,t}^*(p) = (T - h_{1t}) \frac{\varphi_{1,t}(p)}{\alpha_1 + \varphi_{1,t}(a) + \varphi_{1,t}(p)} \quad (3)$$

$$\tau_{2,t}^*(p) = (T - h_{2t}) \frac{\varphi_{2,t}(p)}{\alpha_2 + \varphi_{2,t}(a) + \varphi_{2,t}(p)} \quad (4)$$

$$e_t^* = (w_{1t}h_{1t} + w_{2t}h_{2t} + I_t) \frac{\varphi_{3,t}}{\alpha_3 + \varphi_{3,t}} \quad (5)$$

where

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

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

$$\begin{aligned} \eta_{M+1} &= \psi \alpha_4 \\ \eta_M &= \alpha_4 + \beta \delta_{4,M} \eta_{M+1} \\ &\vdots \\ \eta_t &= \alpha_4 + \beta \delta_{4,t} \eta_{t+1} \\ &\vdots \\ \eta_1 &= \alpha_4 + \beta \delta_{4,1} \eta_2. \end{aligned} \quad (6)$$

where η_t represents the period t expected marginal utility of (log) child quality to the household, i.e., $\eta_t = \partial E_{t-1} V_t(S_t) / \partial \ln k_t$. The variable η_t reflects both the present period flow marginal utility of (log) child quality to the household, given by α_4 , and the discounted marginal value of child quality in terms of future utility. The latter value of current child quality depends on the discount factor 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

$$\begin{aligned}\hat{h}_{1t} &= \frac{A_{1t} - A_{2t}B_{1t}}{1 - A_{2t}B_{2t}} \\ \hat{h}_{2t} &= \frac{B_{1t} - B_{2t}A_{1t}}{1 - A_{2t}B_{2t}},\end{aligned}\tag{7}$$

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 equal to or less than zero, 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 (7), 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

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

Using these optimal labor supply choices, the investment decisions are determined using (1), (2), (3), (4), and (5) after substituting h_{1t}^* and h_{2t}^* into the functions.

B Solving for J_t

In terms of actually solving the constrained optimization problem that defines J_t , we consider the maximization of

$$\alpha_1 \ln(l_{1,t}) + \alpha_2 \ln l_{2,t} + \alpha_3 \ln c_t + \alpha_4 \ln k_t \\ + \lambda_t \{ R_t \tau_{1,t}(a)^{\delta_{1,t}(a)} \tau_{2,t}(a)^{\delta_{2,t}(a)} \tau_{1,t}(p)^{\delta_{1,t}(p)} \tau_{2,t}(p)^{\delta_{2,t}(p)} e_t^{\delta_{3,t}} k_t^{\delta_{4,t}} - k_{t+1} \},$$

where k_{t+1} is the target level for period $t + 1$ child cognitive ability. The constraints are

$$\begin{aligned} l_{1,t} &= T - \tau_{1,t}(a) - \tau_{1,t}(p) - h_{1,t} \\ l_{2,t} &= T - \tau_{2,t}(a) - \tau_{2,t}(p) - h_{2,t} \\ c_t &= w_{1,t}h_{1,t} + w_{2,t}h_{2,t} + I_t - e_t. \end{aligned}$$

Conditional on values of $(h_{1,t}, h_{2,t})$, we can rewrite the first order conditions (FOCs) as

$$\begin{aligned} \frac{\alpha_1 \tau_{1,t}^*(a)}{l_{1,t}^* \delta_{1,t}(a)} &= \lambda_t k_{t+1} \\ \frac{\alpha_1 \tau_{1,t}^*(p)}{l_{1,t}^* \delta_{1,t}(p)} &= \lambda_t k_{t+1} \\ \frac{\alpha_2 \tau_{2,t}^*(a)}{l_{2,t}^* \delta_{2,t}(a)} &= \lambda_t k_{t+1} \\ \frac{\alpha_2 \tau_{2,t}^*(p)}{l_{2,t}^* \delta_{2,t}(p)} &= \lambda_t k_{t+1} \\ \frac{\alpha_3 e_t^*}{c_t^* \delta_{3,t}} &= \lambda_t k_{t+1}. \end{aligned}$$

Conditional on a value of e_t , the FOCs imply

$$\begin{aligned} \tau_{1,t}^*(p)(e_t) &= \frac{\delta_{1,t}(p)}{\delta_{1,t}(a)} \tau_{1,t}^*(a)(e_t) \\ \tau_{1,t}^*(a)(e_t) &= \frac{\delta_{1,t}(a) \alpha_3 (T - h_{1,t}) e_t}{\alpha_1 \delta_{3,t} (M_t - e_t) + \delta_{1,t}(a) \alpha_3 (1 + \frac{\delta_{1,t}(p)}{\delta_{1,t}(a)}) e_t} \\ \tau_{2,t}^*(p)(e_t) &= \frac{\delta_{2,t}(p)}{\delta_{2,t}(a)} \tau_{2,t}^*(a) \\ \tau_{2,t}^*(a)(e_t) &= \frac{\delta_{2,t}(a) \alpha_3 (T - h_{2,t}) e_t}{\alpha_2 \delta_{3,t} (M_t - e_t) + \delta_{2,t}(a) \alpha_3 (1 + \frac{\delta_{2,t}(p)}{\delta_{2,t}(a)}) e_t}, \end{aligned}$$

where total income $M_t = w_{1,t}h_{1,t} + w_{2,t}h_{2,t} + I_t$.

Now since $k_{t+1} = R_t \tau_{1,t}(a)^{\delta_{1,t}(a)} \tau_{2,t}(a)^{\delta_{2,t}(a)} \tau_{1,t}(p)^{\delta_{1,t}(p)} \tau_{2,t}(p)^{\delta_{2,t}(p)} e_t^{\delta_{3,t}} k_t^{\delta_{4,t}}$, given all time inputs, the required amount of expenditures to produce k_{t+1} is


$$\begin{aligned} \hat{e}_t(k_{t+1}, k_t, \tau_{1,t}(a), \tau_{2,t}(a), \tau_{1,t}(p), \tau_{2,t}(p)) &= k_{t+1}^{1/\delta_{3,t}} R_t^{-1/\delta_{3,t}} \tau_{1,t}(a)^{-\delta_{1,t}(a)/\delta_{3,t}} \tau_{2,t}(a)^{-\delta_{2,t}(a)/\delta_{3,t}} \\ &\quad \times \tau_{1,t}(p)^{-\delta_{1,t}(p)/\delta_{3,t}} \tau_{2,t}(p)^{-\delta_{2,t}(p)/\delta_{3,t}} k_t^{-\delta_{4,t}/\delta_{3,t}}. \end{aligned}$$

We solve this problem by substituting for the conditional (on e_t) time supplies of the parents in the right hand side, and finding the solution to

$$\begin{aligned} 0 &= e_t^* - k_{t+1}^{1/\delta_{3,t}} R_t^{-1/\delta_{3,t}} \tau_{1,t}^*(a) (e_t^*)^{-\delta_{1,t}(a)/\delta_{3,t}} \tau_{2,t}(a) (e_t^*)^{-\delta_{2,t}(a)/\delta_{3,t}} \\ &\quad \times \tau_{1,t}(p) (e_t^*)^{-\delta_{1,t}(p)/\delta_{3,t}} \tau_{2,t}(p) (e_t^*)^{-\delta_{2,t}(p)/\delta_{3,t}} k_t^{-\delta_{4,t}/\delta_{3,t}}. \end{aligned}$$


The solution to this problem, e_t^* , does not necessarily satisfy the budget constraint that $e_t^* \in (0, M_t)$. When the solution lies outside this interval, this implies that k_{t+1} cannot be produced given the production technology and the state variables in period t .

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




Notes: This graph shows estimated parameters by child age (see DFW 2014).

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




Notes: This graphs estimated parameters by child age (see DFW 2014).

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 at $t + 1$ from baseline as a result of the transfer.

Figure 4: 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 5: Gain in Average Child Quality from a Restricted Income Transfers (by Age of Transfer)




Notes: A restricted transfer is a transfer φ_R 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.

Figure 6: Distribution of Gains in Child Quality from a Restricted Transfer




Notes: A restricted transfer is a transfer of φ_R 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 7: 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 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 8: Average Gain in Child Quality of Conditional Transfer Program



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

Figure 9: 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 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 ($t = 10$). A conditional transfer program transfers ϕ_G dollars to a household in period $t+1$ if the household meets the target of increasing child quality by $k_{t+1}/k_t \geq \rho_G$. The conditional policy here is the minimum cost combination of policy target and reward to achieve a given improvement in average child quality.

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 2014).

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)	(2)	(3)	(4)
	No Transfer	100 Transfer	250 Transfer	500 Transfer
Child Quality (Latent)	4.15	4.19	4.25	4.35
Mother's Labor Hours	33.5	31.5	28.7	24.3
Father's Labor Hours	42.9	40.5	37.1	31.7
Mother's Active Time	12.5	12.9	13.4	14.2
Father's Active Time	8.29	8.57	8.97	9.63
Mother's Passive Time	10.8	11.1	11.6	12.3
Father's Passive Time	7.69	7.95	8.32	8.94
Mother's Leisure Time	55.2	56.5	58.4	61.2
Father's Leisure Time	53.1	55	57.6	61.7
Child Expenditures	261	268	279	297
Household Consumption	1264	1299	1353	1447

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 $t = 10$. Child quality is the latent level of child quality produced at the end of age 10 and is 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)	(2)	(3)	(4)
	No Transfer	100 Transfer	250 Transfer	500 Transfer
Child Quality (Latent)	4.15	4.23	4.39	4.61
Mother's Labor Hours	33.5	31.7	30.2	29.1
Father's Labor Hours	42.9	40.9	39.5	38.5
Mother's Active Time	12.5	12.9	13.2	13.6
Father's Active Time	8.29	8.56	8.82	9.03
Mother's Passive Time	10.8	11.1	11.4	11.7
Father's Passive Time	7.69	7.94	8.19	8.38
Mother's Leisure Time	55.2	56.3	57.2	57.7
Father's Leisure Time	53.1	54.6	55.5	56.1
Child Expenditures	261	280	357	553
Household Consumption	1264	1293	1317	1334

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 $t = 10$. Child quality is the latent level of child quality produced at the end of age 10 and is 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	.634	.834	.881	.894	.917
1.05	.47	.706	.794	.845	.889
1.1	.283	.54	.698	.766	.855

Notes: Take-up rate is the fraction of the population who either already meet the threshold target at baseline or the households who would be willing to meet the threshold given the reward incentive. The columns measure the reward incentive and the rows measure the performance target.

Table 6: 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	16.7	44	89.4	183
1.05	0	14.1	39.7	84.5	178
1.1	0	10.8	34.9	76.6	171

Notes: Per capita cost is the fraction of the population who would meet the threshold (from table 5) multiplied by the reward level ϕ_G . 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.