

Collegio Carlo Alberto



Household Choices and Child Development

Daniela Del Boca
Christopher Flinn
Matthew Wiswall

No. 149
August 2010

Carlo Alberto Notebooks

www.carloalberto.org/working_papers

© 2010 by Daniela Del Boca, Christopher Flinn and Matthew Wiswall. Any opinions expressed here are those of the authors and not those of the Collegio Carlo Alberto.

Household Choices and Child Development*

Daniela Del Boca
University of Torino
Collegio Carlo Alberto

Christopher Flinn
New York University
Collegio Carlo Alberto

Matthew Wiswall
New York University

August 2010

JEL Classification J13, D1

Keywords: Time Allocation; Child Development, Household Labor Supply.

Abstract

The growth in labor market participation among women with young children has raised concerns about the potential negative impact of the mother's absence from home on child outcomes. Recent data show that mother's time spent with children has declined in the last decade, while the indicators of children's cognitive and noncognitive outcomes have worsened. The objective of our research is to estimate a model of the cognitive development process of children nested within an otherwise standard model of household life cycle behavior. The model generates endogenous dynamic interrelationships between the child quality and employment processes in the household, which are found to be consistent with patterns observed in the data. The estimated model is used to explore the effects of schooling subsidies and employment restrictions on household welfare and child development.

*Part of this research has been supported by a grant from the Institute of Human Development and Social Change at New York University and by a grant from the Collegio Carlo Alberto. Flinn thanks the National Science Foundation and Flinn and Wiswall acknowledge support from the C.V.Starr Center for Applied Economics at New York University as well. For helpful comments, we thank Thomas Demuynck and Robert Sauer and seminar participants at Universidad Pablo de Olavide, Autònoma (Barcelona), Bocconi, the Tinbergen Institute, Tilburg, Collegio Carlo Alberto, CASSR (Sociology, NYU) and conference participants at Collegio Carlo Alberto, Cergy, IFS, EALE/SOLE 2010, and the Etta Chiuri Memorial conference held in Bari. We are responsible for all errors, interpretations, and omissions.

1 Introduction

Economic theory does not provide unambiguous predictions regarding the impact of parental employment on the welfare of children. While family income is necessary to provide for the public and private consumption of the household and its members, as well as for investments undertaken collectively and individually, there are opportunity costs associated with time supplied to the labor market beyond the foregone leisure of the parents. As discussed below, numerous studies have found a positive relationship between the time spent with parents and the quality of a child's outcomes, measured on any of one of several dimensions. The nature of the tradeoff facing household decision-makers is apparent if we think of the household as an enterprise with a complicated set of preferences, outputs, and constraints, as in the seminal work of Becker (1981). From a societal perspective, perhaps the most important output of the household is the number and the "quality" of children it produces. It is helpful to think of there being a production technology for child quality, in which some initial endowment of child quality at birth is augmented during the development process by inputs of time contributed by parents, siblings, other relatives, paid child-care workers, teachers, etc., and by various types of goods in the market, such as formal schooling, toys, books, sporting goods, etc. Given the household's objectives, its mode of decision-making, and the constraints it is confronted with over time, it makes a sequence of time allocation, consumption, and investment decisions at each stage of the child development process. Ignoring the possibility of borrowing and saving for the moment, the more income the household has at a point in time, the more can be spent on child investment goods, among other things. By the same token, decreasing time with the child, holding other inputs fixed, leads to worse child quality outcomes. How does the household properly balance these trade-offs?

It may be useful to compare and contrast the situation that prevails in the standard profit-maximizing firm. In this case, firm behavior is decomposed into a (1) cost-minimization problem, which yields a function that provides the minimum cost required to produce any level of output, and (2) a profit-maximization problem in which the cost function is considered as being predetermined. The solution of these two sequential problems yields the profit-maximizing output quantity and the demand for the factors of production. This is all reasonably straightforward.

In the case of child quality, we face a number of additional challenges, some of them nicely described in Todd and Wolpin (2004). The production process is truly dynamic in the case of child quality, with input decisions from past years typically influencing current investment choices and child quality levels. Unlike the case of the firm, whose objective is to choose output levels so as to maximize profits, the objectives of the parents are not defined only with respect to the quality, or the quantity, of their children. Their utility levels also are determined by their leisure, their consumption of adult private goods, and public goods that affect the welfare of all members of the household. In other words, the household's objective is *not* to maximize child quality given the resources available to them.

Our focus is on the estimation of child outcome technologies, or production functions, that include as arguments a limited number of (potentially observable) factors of production, as well as functions characterizing the dynamic evolution of the budget constraint of the household. Given the arguments we have just presented, it appears necessary to model the household decision-making process in a relatively complete way if we are to disentangle the impact of household preferences, technologies, and constraints on child quality outcomes. Our model utilizes a reasonably standard life cycle framework, in which parents face partially endogenous constraint sets that evolve over time. The Child Development Supplements of the Panel Study of Income Dynamics, hereafter referred to as the PSID-CDS, gives us a substantial amount of useful information, but only at one or two points in time. The keys to being able to use such limited (in a dynamic sense) information to estimate a growth model are (1) assuming age-invariance or parametric age-dependence of the functions and processes describing the households' objectives and constraints and (2) the use of simulation-based estimation techniques that allow us to "fill in" the large numbers of gaps we face in our data on the development process at the household level.

Even under the restrictive assumptions made for purposes of tractability and for parameter identification, we find that we are able to fit many of the patterns observed in the empirical household income, labor supply, child investment, and child outcome processes reasonably well. Using the parameter estimates, we are able to analyze the impact of changes in the time inputs of mothers and fathers on the child development process. Of course, both of these processes are endogenous within the model, so that any changes in the relationship between them must be produced by changes in the parental wage and nonlabor income processes and the prices of the consumption and investment goods. The model is able to generate a number of behavioral links between the child quality and employment processes in the household, which we believe shed some light on dynamic relationships observed in the raw data and the possible impacts of labor market shocks to the parents on the welfare of their children.

There exists a large empirical literature on the relationship between household characteristics, parental employment patterns, and child outcomes, and in section 2 we provide a brief survey. In section 3 we present the model, and section 4 contains a discussion of estimation issues, the data utilized in the empirical work, and some descriptive empirical results. Section 5 contains the model estimates, and some "empirical" comparative statics exercises. In section 6 we conduct two highly stylized policy experiments, primarily in order to illustrate the potential usefulness of the modeling approach taken in the paper to addressing policy issues. Section 7 concludes.

2 Previous Research on the Determinants of Child Development

There is an extensive literature in economics on parental and public investment in children and children's outcomes. Recent surveys have shown that children's cognitive and non-cognitive outcomes are largely determined early in life (Carneiro and Heckman 2003, Ermisch and Francesconi 2005). Inputs applied by families as well as other environmental factors during the early childhood play a very significant role in later cognitive, social, and behavioral outcomes. Several papers have examined this relationship. Todd and Wolpin (2007) estimate a dynamic child quality production function that views child development as a cumulative process, with the final child quality level being determined by the sequence of family and school inputs supplied during the developmental process and on heritable endowments (i.e., initial conditions). Their estimating framework allows for unobserved endowment effects, potentially endogenous input choices, and for the cumulative effects of child investments. Their results show that both contemporaneous and lagged inputs matter in the production of current achievement, and that it is important to allow for unobserved child-specific endowment effects and the endogeneity of inputs. Cunha and Heckman (2008) and Cunha, Heckman, and Schenach (2010) estimate a dynamic factor model of child cognitive and non-cognitive outcomes in a model of skill formation, taking into account the problem of endogeneity of inputs and the unobserved nature of both the inputs and outputs. They find that early environments play a large role in shaping later outcomes, and children's cognitive and non-cognitive outcomes are largely determined early in life.

Our research builds on these previous studies by jointly estimating the production technology of child development within an explicit model of household choices. This strategy accomplishes the goal of "correcting" for the endogeneity of inputs in the estimation of the production technology, as in the work by Todd and Wolpin and Cunha et al., but also allows us to estimate the household preferences that lead to these input decisions, albeit with explicit assumptions about the model of the household. This allows us to conduct more realistic policy experiments by manipulating the time and budget constraints the household faces (e.g. income transfers) in order to understand how households adjust their input choices to changes in the policy environment and how this ultimately affects the child development process. A limitation of our approach relative to Cunha et al is that we consider only household investments in cognitive development, as in the work of Todd and Wolpin, rather than both non-cognitive and cognitive development as in Cunha et al. (2010) However, in contrast to the work of Cunha et al, which considers only a single child investment good, households in our model make a number of specific input choices, ranging from different time inputs to child good expenditures, each with a different time varying productivity. Our model of the household then allows us to incorporate directly a rich variety of household level data, including parental labor supply, wages, and non-labor

income, and relate these data to the child development process. This approach is related to the work of Bernal (2008). She estimates a dynamic model of mothers' choices to control for potential biases that may arise as a result of the fact that women who work and use child care may be systematically different from women who do not, as well as allowing for the child's initial cognitive ability to influence the mother's work decisions. In contrast to Bernal, our work involves a richer set of household choices, including continuous labor supply, time, and expenditure choices of both mothers and fathers.

A large number of studies have assessed the effect of parental time on children's cognitive development. Most studies have used parents' employment as a proxy for time with the children. These studies report evidence that, on the negative side, the loss of parental time with the children has a negative impact on certain measures of the child's well-being (e.g. socio-emotional adjustment and cognitive outcomes), while, on the positive side, the additional labor income has positive implications for expenditures on goods consumed by the child (Brooks-Gunn, Han and Waldfogel, 2001; Ermisch and Francesconi, 2005; Bernal, 2008). It remains unclear which effect is predominant, since the existing literature provides conflicting conclusions. There is wide variation in reported empirical estimates, even for studies based on the same data set. Estimates range from parental employment being detrimental (Baydar and Brooks-Gunn, 1991; Desai et. al., 1989), to its having no effect (Blau and Grossberg, 1992), to its being beneficial (Vandell and Ramanan, 1992). Reasons for the diversity of these results may include the wide range of specifications that are estimated, as well as the common limitation of failing to control for potential biases that may arise due to the endogeneity of parental time and other child quality inputs included in the analyses.

Most studies limit their attention to mother's inputs. The literature on the effects of maternal time on child cognitive outcomes is extensive. Some studies have focused on the timing of the process showing that there are deleterious impacts of maternal employment during the child's first year, while the influence of maternal employment on child outcomes after the first year is ambiguous (Baydar and Brooks-Gunn, 1991; Ruhm, 2004). Many studies have demonstrated that the influence of maternal employment on child outcomes differs by the economic and demographic characteristics of mothers and their families, suggesting that maternal employment may be more harmful for children from advantaged backgrounds: children from wealthier families, non-Hispanic white children, and children from intact families. Desai, Chase-Lansdale and Michael (1989) find that maternal employment negatively influences children from higher income families, but not children from middle or low income families, while Waldfogel, Han and Brooks-Gunn (2002) find a persistent negative effect of mother's employment on cognitive test scores for non-Hispanic white children, but not for African-American or Hispanic children.

Very few studies have used direct measures of the time parents spend with children to examine the relationship between parental investments and children's cognitive development. Time diary data suggest that women's entry into the labor force is associated with changes in time use but employment is not a perfect proxy of time with children. Huston

and Aronson (2005) find that a measure of the mother's time with children even relates negatively to child language skills, though there are obvious reverse causality issues that may be important. Their studies are limited to the investigation of behaviors and outcomes in the first two years of life and do not take into account differences in the initial endowments of children. More recently, Hsin (2008), using the Child Development Supplement of the PSID, investigated the effect of maternal time with children during pre-school years on the child's cognitive outcomes, while conditioning on characteristics of children that may lead to bias the estimated impact of maternal time allocations. For example, children clearly differ in their initial endowments such as innate cognitive ability, health status, and their physical development. Mothers may respond to observed difficulties faced by children by spending more time with them. She finds a positive and persistent effect of the time mothers spend with children on children's language development, but only among children who spend time with verbally-skilled mothers.

While mother's time is a crucial input in the production process of child outcomes, father's time may be equally productive, especially in some stages in the child's development process; time spent with children by fathers has increased over time, partly offsetting the decline in mother's time spent with the child. Studies considering father's time show that fathers' care for infants is no better or worse than other types of arrangements (Averett et al (2005), and the amount of time a father spends with children is affected by the gender composition of the children (Lundberg et al 2006). Several studies suggest that there is no long term benefit of paternal investments on children's achievement and behavior (Yeung, Hill, and Duncan, 1999; Haveman and Wolfe, 1995).

In addition to the literature on parental time, another literature examines whether family income influences child development (see surveys by Mayer 1997, and Almond and Currie 2010). Despite a large body of evidence documenting the association between family income and child development, there is much controversy about whether these correlations can be given causal interpretations. Like the parental time literature, the results in the family income literature are mixed. Using instrumental variable strategies, the estimates reported in Dahl and Lochner (2008) suggest some positive effects of family income on children's (short-run) outcomes and Loken, while Loken (2010) finds little, if any, impact of family income. Using a sibling differences in family income to net out permanent differences in family environment, Levy and Duncan (2000) find that family income is important for children's educational attainment, whereas Blau (1999) finds a small effect of family income on child outcome. Loken, Mogstad, and Wiswall (2010) show the importance of model specification as restricted linear models, as estimated previously, may show very small effects, while non-linear models reveal a concave relationship with stronger effects of family income on child outcomes for poor families and little effect for richer families, as is consistent with the theoretical literature (Becker and Tomes 1979).

An important contribution of our research is that we can show the different connections between the level of household income and childhood development. A higher level of family income does not necessarily indicate a higher level of family resources being provided to

children. This is because for most households, most income is derived from labor market participation, and labor supply directly trades off with time spent with children. Parents with higher levels of income may then be working more in the labor market and providing lower levels of time investments in their children. This channel may dampen or even reverse the assumed positive relationship between income and child development. Our model allows us to examine each of these channels through which family resources and household decision making affect the production of child outcomes.

3 Model

This section develops the model that is the basis of our empirical estimation. As noted in the Introduction, 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 forms of the life-cycle demand functions that allows us to include a relatively large number of endogenous variables in a straightforward way. The special characteristics of the decision rules also allow us to sort out identification issues when we discuss estimation issues in the following section.

3.1 Timing and Preferences

The model begins with the birth of a child. The household makes decisions in each period of a child’s life, where the child’s age is indexed by t . We consider only one child families and leave for future work extensions of this modeling framework to multiple child households. Parents make investments in child quality from the first period of the child’s life, $t = 1$, through the last developmental period, T . At this “terminal” point (from the perspective of the parents’ investment in the child), the child has reached adulthood and adult outcomes depend (in part) on the level of child quality obtained at this point.¹

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

¹The terminal date T need 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.

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

where $\sum_j \alpha_j = 1$.

3.2 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:

$$k_{t+1} = f(k_t, \tau_{1t}, \tau_{2t}, z_{1t}, z_{2t}, e_t) = R_t \tau_{1t}^{\delta_{1t}} \tau_{2t}^{\delta_{2t}} z_{1t}^{\delta_{3t}} z_{2t}^{\delta_{4t}} e_t^{\delta_{5t}} k_t^{\delta_{6t}} \quad (2)$$

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

While the Cobb-Douglas form restricts the substitution possibilities, we allow the productivities of the various inputs to vary over the age of the child. This allows us to capture the important insights in the economics and child development literatures 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. An extension of the model is to assume there are stochastic shocks to child quality production, which may be attractive for purposes of estimation, though under our functional form assumptions the presence of such a shock has no substantive impact on the household's decision rules. We will return to this topic below.

3.3 Dynamic Problem

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

$$V_t(S_t) = \max_{l_{1t}, \tau_{1t}, l_{2t}, \tau_{2t}, z_{1t}, z_{2t}, e_t} u(l_{1t}, l_{2t}, c_t, k_t) + \beta E_t V_{t+1}(S_{t+1}), \quad (3)$$

$$\text{s.t. } TT = l_{jt} + h_{jt} + \tau_{jt} + z_{jt}, \quad j = 1, 2 \quad (4)$$

$$c_t + e_t = w_{1t} h_{1t} + w_{2t} h_{2t} + I_t \quad (5)$$

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

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

β ($\in [0, 1)$) is the discount factor, and E_t denotes the conditional expectation operator with respect to the period t information set. The state variable vector at the birth of the child are the initial conditions of the problem, $S_1 = (k_1 \ w_{11} \ w_{21} \ I_1)$.

The constraint set faced by the household in period t consists of time and market good expenditures restrictions. We assume that each parent has a time endowment of TT 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.

3.4 Terminal Value

Parental investments in child quality are limited to the first T period's of the child's life. The terminal level of child quality is then k_{T+1} . Beyond T , the parents make only labor supply and consumption decisions, since $k_t = k_{T+1}$ for $t = T + 1, T + 2, \dots$. Then we can write the households optimization problem at time $T + 1$ as

$$V_{T+1}(w_{1,T+1} \ w_{2,T+1} \ I_{T+1}; k_{T+1}) = \tilde{V}_{T+1}(w_{1,T+1}, w_{2,T+1}, I_{T+1}) + \psi\alpha_4 \ln k_{T+1},$$

where

$$\begin{aligned} \tilde{V}_{T+1}(w_{1,T+1}, w_{2,T+1}, I_{T+1}) = \\ \max_{l_{1,T+1}, l_{2,T+1}} \quad & \alpha_1 \ln l_{1,T+1} + \alpha_2 \ln l_{2,T+1} + \alpha_3 \ln(w_{1,T+1}(TT - l_{1,T+1}) + w_{2,T+1}(TT - l_{2,T+1}) + I_{T+1}) \\ & + \beta E_{t+1} \tilde{V}_{T+2}(w_{1,T+2}, w_{2,T+2}, I_{T+2}). \end{aligned}$$

A few things to note about this expression are the following. First, assuming that the household is infinitely-lived,² the household problem after period T becomes stationary, so that we can write

$$\tilde{V}_{T+s}(\tilde{S}_{T+s}) = \tilde{V}(\tilde{S}_{T+s}), \quad s = 1, 2, \dots,$$

where $\tilde{S}_t = (w_{1t} \ w_{2t} \ I_t)$. Second, we have assumed that the “terminal” value of child quality in period $T + 1$ is given by $\psi\alpha_4 \ln k_{T+1}$. There are at least two ways to view this component of the household's payoff function.

1. If we made the stringent assumption that child quality was fixed at k_{T+1} for the remainder of the child's life (or at least, the household's view of the child's quality was fixed at k_{T+1} “forever), then the utility yield to the infinitely-lived household would be

$$\alpha_4 \ln k_{T+1} + \beta\alpha_4 \ln k_{T+1} + \beta^2\alpha_4 \ln k_{T+1} + \dots = (1 - \beta)^{-1}\alpha_4 \ln k_{T+1},$$

in which case $\psi = (1 - \beta)^{-1}$.

²It is not strictly necessary that the household be infinitely-lived. All of the properties discussed follow if the has a constant probability of death, say π , the value of which is subsumed within the discount factor β .

2. Somewhat more realistically, we can think of the quality level k_{T+1} as being an input, or serving as an initial condition, into a subsequent development phase of the child, one in which the household makes no direct money or time investments. We assume that the expected value of this stream of future quality levels can be represented by $\psi \ln k_{T+1}$.

Finally, it is apparent that with no savings or borrowing possibilities and with no state dependence in wages and/or nonlabor income (i.e., the wage and nonlabor income processes are strictly exogenous), once the child investment process is finished, the household solves a sequence of static optimization problems. Thus the period s optimization problem, $s = T + 1, T + 2, \dots$, reduces to

$$\max_{l_{1,s}, l_{2,s}} \alpha_1 \ln l_{1,s} + \alpha_2 \ln l_{2,s} + \alpha_3 \ln(w_{1,s}(TT - l_{1,s}) + w_{2,s}(TT - l_{2,s}) + I_s).$$

All of this implies that we can write the period T optimization problem as

$$\begin{aligned} V_T(w_{1T}, w_{2T}, I_T, k_T) = & \\ & \max_{l_{1T}, \tau_{1T}, l_{2T}, \tau_{2T}, e_T} \alpha_1 \ln l_{1T} + \alpha_2 \ln l_{2T} + \alpha_3 \ln e_T + \alpha_4 \ln k_T \\ & + \psi \alpha_4 \{ \delta_{1T} \ln \tau_{1T} + \delta_{2T} \ln \tau_{2T} + \delta_{3T} \ln z_{1T} + \delta_{4T} \ln z_{2T} + \delta_{5T} \ln e_T + \delta_{6T} \ln k_T \} \\ & + \beta E_T \tilde{V}(\tilde{S}_{T+1}) \end{aligned}$$

From the point of view of the choice problem, we can write this expression as

$$\begin{aligned} V_T(w_{1T}, w_{2T}, I_T, k_T) = & \max_{l_{1T}, \tau_{1T}, l_{2T}, \tau_{2T}, z_{1T}, z_{2T}, e_T} \alpha_1 \ln l_{1T} + \alpha_2 \ln l_{2T} + \alpha_3 \ln e_T \\ & + \psi \alpha_4 \{ \delta_{1T} \ln \tau_{1T} + \delta_{2T} \ln \tau_{2T} + \delta_{3T} \ln z_{1T} + \delta_{4T} \ln z_{2T} + \delta_{5T} \ln e_T \} \\ & + Q_T(k_T, \tilde{S}_{T+1}), \end{aligned} \quad (6)$$

where $Q_T(k_T, \tilde{S}_{T+1})$ is not a function of any current period choices. Thus the time allocation and consumption decisions are determined solely from the expression on the first two lines of the right hand side of (6).

3.5 Model Solution

We devote some time to describing the solution to the model, which will be important in understanding the ability of the model to fit the data and, more formally, in evaluating the ability of our proposed estimator to recover the primitive parameters that characterize the model.

As is clear from the nature of the production technology, there are never any corner solutions to the household input choice problem during the investment period.³ However,

³If 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 $\alpha_4 > 0$.

we do allow for corner solutions in labor supply as labor supply for either or both parents may be 0 in a given period. We can write the conditional factor demands for child inputs, where we are conditioning on labor supply choices and nonlabor income, as

$$\tau_{1t} = (TT - h_{1t}) \frac{\varphi_{1t}}{\alpha_1 + \varphi_{1t} + \varphi_{3t}} \quad (7)$$

$$\tau_{2t} = (TT - h_{2t}) \frac{\varphi_{2t}}{\alpha_2 + \varphi_{2t} + \varphi_{4t}} \quad (8)$$

$$z_{1t} = (TT - h_{1t}) \frac{\varphi_{3t}}{\alpha_1 + \varphi_{1t} + \varphi_{3t}} \quad (9)$$

$$z_{2t} = (TT - h_{2t}) \frac{\varphi_{4t}}{\alpha_2 + \varphi_{2t} + \varphi_{4t}} \quad (10)$$

$$e_t = (w_{1t}h_{1t} + w_{2t}h_{2t} + I_t) \frac{\varphi_{5t}}{\alpha_3 + \varphi_{5t}} \quad (11)$$

where

$$\varphi_{jt} = \beta \delta_{jt} \eta_{t+1}, \quad j = 1, \dots, 5.$$

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

$$\begin{aligned} \eta_{T+1} &= \psi \alpha_4 \\ \eta_T &= \alpha_4 + \beta \delta_{6,T} \eta_{T+1} \\ &\vdots \\ \eta_t &= \alpha_4 + \beta \delta_{6,t} \eta_{t+1} \\ &\vdots \\ \eta_1 &= \alpha_4 + \beta \delta_{6,2} \eta_2. \end{aligned}$$

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

The solution to the spousal labor supplies problem in period t also has a simple form. Define two ‘‘latent’’ labor supply variables in period t by

$$\begin{aligned} \hat{h}_{1t} &= \frac{A_{1t} - A_{2t}B_{1t}}{1 - A_{2t}B_{2t}} \\ \hat{h}_{2t} &= \frac{B_{1t} - B_{2t}A_{1t}}{1 - A_{2t}B_{2t}}, \end{aligned} \quad (12)$$

where

$$\begin{aligned}
A_{1t} &= \frac{w_{1t}TT(\alpha_3 + \varphi_{5t}) - (\alpha_1 + \varphi_{1t} + \varphi_{3t})I_t}{w_{1t}(\alpha_1 + \alpha_3 + \varphi_{1t} + \varphi_{3t} + \varphi_{5t})} \\
A_{2t} &= \frac{w_{2t}(\alpha_1 + \varphi_{1t} + \varphi_{3t})}{w_{1t}(\alpha_1 + \alpha_3 + \varphi_{1t} + \varphi_{3t} + \varphi_{5t})} \\
B_{1t} &= \frac{w_{2t}TT(\alpha_3 + \varphi_{5t}) - (\alpha_2 + \varphi_{2t} + \varphi_{4t})I_t}{w_{2t}(\alpha_2 + \alpha_3 + \varphi_{2t} + \varphi_{4t} + \varphi_{5t})} \\
B_{2t} &= \frac{w_{1t}(\alpha_2 + \varphi_{2t} + \varphi_{4t})}{w_{2t}(\alpha_2 + \alpha_3 + \varphi_{2t} + \varphi_{4t} + \varphi_{5t})},
\end{aligned}$$

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. The solution to the optimization problem in terms of the “latent” labor supply choices is derived from functions similar to reaction functions, which express the optimal latent labor supply choice of parent i as a linear function of the latent labor supply choice of parent i' , $i' \neq i$. These functions are

$$\begin{aligned}
\hat{h}_{1t} &= A_{1t} - A_{2t}\hat{h}_{2t} \\
\hat{h}_{2t} &= B_{1t} - B_{2t}\hat{h}_{1t}.
\end{aligned}$$

If the latent labor supplies on the right hand sides are set to zero, it is apparent that the condition required for the conditional latent labor supplies to both be 0 is

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

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

Going back to the “full” solutions to the model given in (12), 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 (7), (8), (9), (10), and (11) after substituting h_{1t}^* and h_{2t}^* into the functions.

3.6 Characteristics of Decision Rules

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

The model as it has been developed so far does not include “shocks” or uncertainty in any form. Once again, due to the properties of the decision rules, such “realistic” changes to the model setup can be easily accommodated. For example, say that the child quality production function is altered to include a stochastic total factor productivity, that is,

$$k_{t+1} = f_t(R_t, \tau_{1t}, \tau_{2t}, z_{1t}, z_{2t}, e_t, k_t) = R_t \tau_{1t}^{\delta_{1t}} \tau_{2t}^{\delta_{2t}} z_{1t}^{\delta_{3t}} z_{2t}^{\delta_{4t}} e_t^{\delta_{5t}} k_t^{\delta_{6t}},$$

where R_t has a nondegenerate distribution H_t . Then we can write

$$\ln k_{t+1} = \overline{\ln R_t} + \delta_{1t} \ln \tau_{1t} + \delta_{2t} \ln \tau_{2t} + \delta_{3t} \ln z_{1t} + \delta_{4t} \ln z_{2t} + \delta_{5t} \ln e_t + \delta_{6t} \ln k_t + \varepsilon_t, \quad (13)$$

where $\varepsilon_t \equiv \ln R_t - \overline{\ln R_t}$ and $E_t \ln R_t = \overline{\ln R_t}$, assuming that this expectation exists. No other restrictions are necessary on the distribution of the sequence $\{\ln R_t\}_{t=1}^T$, since the decision rules are independent of the $\{\ln k_t\}_{t=1}^T$.

It is also not necessary to assume that the household knows future values of the wage and nonlabor income process when making period t decisions *under our assumption of no borrowing or saving*. Because of this assumption, labor supply decisions were made as if the model was static in nature. Due to this, future wages or nonlabor income receipts cannot influence current period decisions, and all decisions are consistent with various assumptions about how households form expectations, including perfect foresight, rational expectations, and adaptive expectations. Given the ability to save and borrow, current period labor supply decisions will be affected by future wage expectations, and the model would have to take a stand on how these were formed.

4 Econometric Issues

We now discuss the manner in which we have taken the model developed above to the data, and we provide some details concerning identification issues. A key element of taking any model to data is the manner in which heterogeneity is included, if it is included in the model at all, and the nature of the randomness in the model that is a requirement for there to exist a well-posed estimation problem.

4.1 Econometric Specification

We begin by discussing some of our assumptions regarding the model specification. As noted above, we allow the production function parameters to vary with the age of the child, but do not allow any further heterogeneity in that function. That is, we assume that all families possess the same child production technology.⁴ In order to economize on parameters, we assume that the parameters $\{\delta_{1t}, \delta_{2t}, \delta_{3t}, \delta_{4t}, \delta_{5t}, \delta_{6t}\}_{t=1}^T$ vary as follows:

$$\delta_{jt} = \exp(\gamma_{j0} + \gamma_{j1}t), \quad j = 1, \dots, 6; \quad t = 1, \dots, T.$$

Similarly, we assume that TFP, R_t , is given by

$$R_t = \exp(\gamma_{00} + \gamma_{01}t + \eta_t),$$

where η_t is i.i.d. $N(0, \sigma_R^2)$. Thus the production process is characterized by 15 parameters, $(\{\gamma_{j0}, \gamma_{j1}\}_{j=0}^6, \sigma_R)$.

Household preferences are assumed to be fixed over time, however, we do allow heterogeneity in utility function preferences across households. We specify the distribution of preferences parameters as $G(\alpha; \theta)$, where G is a parametric distribution function characterized by the finite-dimensional parameter vector θ , and with the three dimensional vector $\alpha = (\alpha_1 \ \alpha_2 \ \alpha_3)'$ defined such that $0 < \alpha_1 + \alpha_2 + \alpha_3 < 1$, $\alpha_j > 0$, $j = 1, 2, 3$. A household objective is determined by a draw from the distribution G , with $\alpha_4 = 1 - \alpha_1 - \alpha_2 - \alpha_3$.

In terms of the specific functional form for G , we assume that it is generated as follows. The 3×1 vector ν is normally distributed with

$$\nu \sim N(\mu_\alpha, \Sigma_\alpha),$$

where μ_α is a 3×1 vector and Σ_α is a 3×3 covariance matrix of full rank. Define $D = 1 + \sum_{j=1}^3 \exp(\nu_j)$. Then a draw ν from the trivariate normal is mapped into α as

$$\begin{aligned} \alpha_1 &= D^{-1} \exp(\nu_1) \\ \alpha_2 &= D^{-1} \exp(\nu_2) \\ \alpha_3 &= D^{-1} \exp(\nu_3) \\ \alpha_4 &= D^{-1} \end{aligned}$$

⁴Below, we do consider another version of the model where the productivity of each parent's time in producing child quality is affected by the parent's level of education.

The c.d.f. of α is given by

$$G(\alpha) = \iiint \chi[D^{-1} \exp(\nu_1) \leq \alpha_1] \chi[D^{-1} \exp(\nu_2) \leq \alpha_2] \chi[D^{-1} \exp(\nu_3) \leq \alpha_3] dF(\nu | \mu_\alpha, \Sigma_\alpha).$$

Thus the population distribution of α is characterized in terms of the parameter vectors μ_α and $vec(\Sigma_\alpha)$, the vectorization of the nonredundant elements in Σ_α .⁵

Next we discuss our specification of the evolution of the wage and nonlabor income processes. We assume that stochastic terms associated with these processes have the following structure:

$$\begin{bmatrix} \ln w_{1,t} \\ \ln w_{2,t} \end{bmatrix} = \begin{bmatrix} \mu_{1,t} \\ \mu_{2,t} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix},$$

where

$$\begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix} \stackrel{i.i.d.}{\sim} N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{bmatrix} \right), \text{ for all } t$$

The terms $\mu_{1,t}$ and $\mu_{2,t}$ are the means of the log wage draws of the mother (1) and father (2) at time t . When the process is estimated, we assume that

$$\ln \mu_{jt} = \mu_j^0 + \mu_j^1 s_j + \mu_j^2 age_{jt} + \mu_j^3 age_{jt}^2, \quad j = 1, 2,$$

where s_j is the completed schooling level (which is time invariant) of parent j .

In terms of the nonlabor income process, there are a large number of households with no nonlabor income in a given period, so we consider this process to be a truncated version of a latent variable process in levels (instead of logs). In particular, let

$$I_t^* = \mu_{3,t} + \varepsilon_{3,t}, \tag{14}$$

be the latent nonlabor income in period t , with a mean given by $\mu_{3,t}$ and where $\varepsilon_{3,t} \stackrel{i.i.d.}{\sim} N(0, \sigma_{33})$, $t = 2, 3, \dots$. The actual nonlabor income process is given by

$$I_t = \max(0, I_t^*), \text{ for all } t. \tag{15}$$

⁵When estimating Σ_α , it is necessary to choose a parameterization that ensures that any estimate $\hat{\Sigma}_\alpha$ is symmetric, positive definite. The most straightforward way of doing so is to use the Cholesky decomposition of Σ_α . There are 10 parameters to estimate, with $\mu_\alpha = [\mu_{\alpha,1}, \mu_{\alpha,2}, \mu_{\alpha,3}]$ and

$$\Sigma_\alpha = \begin{bmatrix} c_1^2 & c_1 c_2 c_{12} & c_1 c_3 c_{13} \\ c_1 c_2 c_{12} & c_2^2 & c_2 c_3 c_{23} \\ c_1 c_3 c_{13} & c_2 c_3 c_{23} & c_3^2 \end{bmatrix},$$

where c_1, c_2, c_3 are constrained to be positive using the exponential function, and c_{12}, c_{13}, c_{23} are constrained to be between -1 and 1 using the logit function. These constraints ensure that the diagonal elements are strictly positive and the matrix has proper scaling for the numerical implementation of the Cholesky decomposition. (see Pinheiro and Bates, 1996).

Since we found little relationship between the observed characteristics of parents and the nonlabor income process, $\mu_{3,t}$ is assumed to be constant across households and over time in the population.

Note that we have assumed that the wage and nonlabor income processes are independently distributed over time. While allowing for dependence in these processes (of the exogenous type⁶) is straightforward and does not complicate the solution of the model, due to the nature of the data available to us we found it impossible to obtain credible estimates of the parameters characterizing dependence. Over the sample period, the PSID is administered every two years, while our model is based on annual time periods. Estimating autocorrelation parameters for a yearly process off of biannual observations leads to a classic aliasing problem. When we allowed the disturbances in the wage equations to follow a first order autoregressive process, point estimates of the autocorrelation parameters were strongly negative, which we found not to be credible. As a result, we have restricted the processes to be conditionally independent over time, with all dependence arising through mean dependence on time-invariant observable heterogeneity.

4.2 Measuring Child Quality

In terms of deriving the mapping between unobserved (latent) child quality, k_t , and measured child quality, k_t^* , we proceed as follows. Consistent with prior research on this subject, we consider child quality to be inherently unobservable to the analyst, though we do assume that it is observable by household members, as it is a determinant of the household utility level and is a (potential) input into the decision-making process. In actuality, most cognitive test scores, such as the one used in the empirical work, are simple sums of the number of questions answered correctly by the test-taker. If the child of age t has a quality level of k_t , we consider the probability that they correctly answer a question of difficulty d to be

$$p(k_t, d).$$

It is natural to assume that p is nondecreasing in its first argument for all d and is nonincreasing in its second argument for all k_t . We will think of the Letter-Word test used in the empirical analysis as consisting of equally “difficult” questions, and will drop the argument d for simplicity.

Given a cognitive ability test consisting of \bar{k} items of equal difficulty, the number of correct answers, k_t^* , is distributed as a Binomial random variable with parameters $(\bar{k}, p(k_t))$. Note that the inherent randomness in the test-taking process ensures that the mapping between k and k^* is stochastic. The measurement process implies that a given child of quality k_t has a stochastic probability of answering some number of the questions correctly, up to the test score ceiling of \bar{k} . Our measurement model then achieves two goals: (i) we

⁶By “nonendogenous,” we mean all dependence that is *not* generated by past household choices, such as previous labor supply decisions.

map a continuous latent child quality defined on $(0, \infty)$ into a discrete test score measure imposing the measurement floor at 0 and ceiling at \bar{k} , and (ii) we allow for the possibility of measurement error so that a child’s score may not perfectly reflect her latent quality. Previous research has often used linear (or log linear) continuous measurement equations (e.g. Cunha and Heckman 2008). Our approach differs from this in using a measurement process that explicitly recognizes the discrete and finite nature of the test score measure.

In order to estimate the model, we do have to take a position on the form of the function $p(k)$. Besides being nondecreasing in k , we would like to have it possess the properties: $\lim_{k \rightarrow 0} p(k) = 0$ and $\lim_{k \rightarrow \infty} p(k) = 1$. We choose the following function that satisfies these restrictions as

$$\begin{aligned} p(k; \lambda) &= \frac{\exp(\lambda_0 + \lambda_1 \ln k)}{1 + \exp(\lambda_0 + \lambda_1 \ln k)} \\ &= \frac{\exp(\lambda_0)k_t^{\lambda_1}}{1 + \exp(\lambda_0)k_t^{\lambda_1}}, \quad \lambda_1 > 0. \end{aligned} \tag{16}$$

Unsurprisingly, we will have to restrict the values of λ_0 and λ_1 to aid in model identification. This is discussed in the following section.

For each child we observe two measures of child quality at two different ages. We use the initial measure of child quality as our initial condition. However, to solve the model, we require an initial level of latent child quality k_t , not the measure k_t^* . We map the initial measure into the initial latent child quality by assuming that, without previous observations on the process prior to our initial measure, we have “total ignorance” regarding a given individual’s value of p . We represent our initial prior as a Beta distribution with parameters $(1, 1)$, or which is uniform on $[0, 1]$. We then observe the test score, allowing us to update our prior and produce a posterior distribution on p , which is also Beta (a conjugate distribution for the Binomial). The posterior distribution for p is then Beta with parameters $(1 + k_t^*, (\bar{k} - k_t^*) + 1)$, where k_t^* is the number of correct answers out of the $\bar{k} = 57$ items.

From the posterior distribution of p , $Beta(1 + k_t^*, (\bar{k} - k_t^*) + 1)$, we draw pseudo random draws of p . Let \tilde{p} be a given draw. We can invert (16) to obtain

$$k_t = \exp(\lambda_1^{-1} \{ \ln(\frac{\tilde{p}}{1 - \tilde{p}}) - \lambda_0 \}).$$

From this initial value of k_t , we then begin the construction of this particular sample path. When we get to the period of the second measurement when the child is age $t' > t$, the test score is drawn from a binomial distribution with parameters $(\bar{k}, p(k_{t'}))$ as described above.

4.3 Identification

In this discussion we indicate the manner in which the behavioral parameters characterizing the model can be recovered from the data at our disposal in a reasonably straightforward manner *given the requisite data*. The estimator we actually implement has several advantages (both theoretical and practical), to be detailed below, over those discussed in this section. But in a reasonably complex dynamic model it is probably useful to develop some intuition as to the key sources of identifying information under our modeling assumptions.

We first consider the estimation of production technology parameters. As noted below, there are 15 parameters to be determined in our parameterization. Estimation of all parameters is hampered by the fact that the PSID-CDS data are not collected on a yearly basis. We employ CDS data that were collected in 1997 and 2002. The timing convention we use is that the child quality measures observed in year t are generated by year t child quality, k_t . Thus, in 2002 for example, for a child of age t in that year, the measured child quality and inputs in 2002 map into the child quality measure of 2003. Unfortunately, that information is unavailable to us. We will carry out our discussion assuming that it is, and then we will consider how our estimation strategy has been adapted to deal with the sizable amount of missing data.

If we did have access to the *actual* level of child quality in successive periods recovery of the δ parameters would be completely straightforward. Say that we had child test score data for both the years 2003 and 2002, as well as the investment information concurrent with the 2002 child quality measure. Using the specification given in (13) and substituting in the functional form of the age-dependent production function parameters, we have

$$\begin{aligned} \ln k_{i,t_i+1} &= \ln R_{i,t_i} + \delta_{1,t_i} \ln \tau_{1,i,t_i} + \delta_{2,t_i} \ln \tau_{2,i,t_i} + \delta_{3,t_i} z_{1t} + \delta_{4,t_i} z_{2t} + \delta_{5,t_i} \ln e_{i,t_i} + \delta_{6,t_i} \ln k_{i,t_i} \\ &= \gamma_{0,0} + \gamma_{0,1} t_i + \exp(\gamma_{1,0} + \gamma_{1,1} t_i) \ln \tau_{1,i,t_i} + \exp(\gamma_{2,0} + \gamma_{2,1} t_i) \ln \tau_{2,i,t_i} \\ &\quad + \exp(\gamma_{3,0} + \gamma_{3,1} t_i) \ln z_{1,i,t_i} + \exp(\gamma_{4,0} + \gamma_{4,1} t_i) \ln z_{2,i,t_i} \\ &\quad + \exp(\gamma_{5,0} + \gamma_{5,1} t_i) \ln e_{i,t_i} + \exp(\gamma_{6,0} + \gamma_{6,1} t_i) \ln k_{i,t_i} + \varepsilon_{i,t_i}, \\ &\equiv X(t_i, \tau_{1,i,t_i}, \tau_{2,i,t_i}, z_{1,i,t_i}, z_{2,i,t_i}, e_{i,t_i}, k_{i,t_i}; \gamma) + \varepsilon_{i,t_i}; \quad i = 1, \dots, N. \end{aligned}$$

where i denotes the household, t_i is the age of the child in household i at the time of the 2002 CDS survey, and $\varepsilon_{i,t_i} \stackrel{i.i.d.}{\sim} N(0, \sigma_R^2)$. Then subject to the usual full rank conditions, the nonlinear least squares estimator

$$\hat{\gamma}_{NLLS} = \arg \min_{\gamma} \sum_{i=1}^N (\ln k_{i,t_i+1} - X(t_i, \tau_{1,i,t_i}, \tau_{2,i,t_i}, z_{1,i,t_i}, z_{2,i,t_i}, e_{i,t_i}, k_{i,t_i}; \gamma))^2. \quad (17)$$

is a consistent estimator of γ , that is, $\text{plim} \hat{\gamma}_{NLLS} = \gamma$. The “full rank conditions” in this case primarily means that not all households choose the same values of investments, which is trivially satisfied in the data, and that not all households have a child of the same age. Since the parameters characterizing the production function are a linear function of age, it

is enough that the sample contain children of two different ages for the full rank condition to be satisfied.⁷ A consistent estimator for the final production process parameter is given by:⁸

$$\hat{\sigma}_R^2 = N^{-1} \sum_{i=1}^N (\ln k_{i,t_i+1} - X(t_i, \tau_{1,i,t_i}, \tau_{2,i,t_i}, z_{1,i,t_i}, z_{2,i,t_i}, e_{i,t_i}, k_{i,t_i}; \hat{\gamma}_{NLLS}))^2.$$

Of course, the argument given above requires that not only are child quality measures available for two successive years, it also assumes that child quality is observed. As discussed above, we only observe a test score measure k^* that allows us to generate a distribution of values of k given our prior distribution and knowledge of the parameters characterizing $p(k)$. If we had access to measures k_{i,t_i+1}^* and k_{i,t_i}^* , then the model and the measurement process imply a mapping between the two distributions that is a function of the observed inputs at time t_i , the measured inputs, and the parameter vector γ . While estimates of γ could not be recovered simply using the NLLS estimator described in (17), a method of moments or maximum likelihood estimator could be defined to recover γ .

In terms of the identification of the wage and income processes, the PSID contains multiple observations of wages (for those parents who work) and non-labor income over several periods across the child development process, including periods for which we do not have child quality data from the CDS. In estimating the nonlabor income process, which is assumed to be independent of the wage processes, there are no issues. The data generating process (DGP) for I_t is described in (14) and (15), under the i.i.d. assumption on the normally distributed disturbances, this is simply a Tobit, and all that is to be estimated is the constant term and the variance of the disturbance, which can be done using a maximum likelihood estimator with cross-sectional data alone.

While we have access to wage observations for multiple periods, wage observations are nonrandomly missing due to the significant number of corner solutions associated with labor supply choices. When one or both parents is not in the labor market, we do not observe the wage. Under our model specification, we can “correct” our estimator of model parameters for the nonrandomly missing data using our the DGP structure from the model. In this case, both the wage processes and the parameters characterizing preferences and production possibilities must be simultaneously estimated.

Each household, at every point in time, makes seven decisions. The time-invariant household utility function is characterized in terms of 3 free parameters, since $\sum_{j=1}^4 \alpha_j = 1$. Other parameters characterizing preferences that remain to be considered are the discount

⁷If the production parameters were a quadratic function of child age, then the full rank condition would require that the sample contain children of at least three different ages, and so on.

⁸Note that if we have observations of child quality and inputs for children of every age, we could consistently estimate the production function parameters δ_{kt} , $k = 1, \dots, 6$, directly using an OLS regression of child quality on inputs for each age. The absence of this information for every child age necessitates some functional form restriction on how these parameters vary by child age.

rate, β , and the parameter ψ that, in conjunction with α_4 , determines the terminal valuation of child quality at this stage of the development process. Both the β and ψ parameters are assumed to be homogeneous in the population.

If we condition on values of β and ψ , then the marginal distribution of α is nonparametrically identified given only one period of observed input demands, labor supplies, and total household income per household. To see this, simply note that by the structure of the production and utility functions, input demands are positive in every period of the production process. The conditional (on labor supply and household income) demand functions are given by (7), (8), (9), (10), and (11). Given β and ψ , this system of equations can be inverted to yield unique values of α_1 , α_2 , and α_3 for each household in the sample. Then the empirical distribution of these values is the nonparametric maximum likelihood estimator of $G(\alpha)$, with the estimator conditional on estimates of the production function parameters and values of β and ψ .

What remains is the determination β and ψ . Under our assumption of homogeneity of these parameters, we could use any household in which both parents work, in conjunction with our estimates of α_1 , α_2 , and α_3 for that household, to determine values of β and ψ . This is accomplished by using the labor supply decisions evaluated at the actual hours choices and wage offers to back these two values out.

While our identification arguments (hopefully) are reasonably convincing at some level, there remain some conceptual problems. Most obvious is the multiplicity of estimates of G , β , and ψ that are available. We have access to two observations on input demands, household labor supplies, and total income for all households in the subsample that we utilize. Conditional on a values of β and ψ , this means that two separate estimates of the household α vector are available. There is no possibility that these “backed out” values of the α vector will be the same at both CDS dates. Of course, one could rationalize input choices at several points in time by adding classical measurement error to the conditional demand equations, but adding measurement error is inherently arbitrary and creates other identification problems.

The main objective of our discussion of identification has been to illustrate that there is a substantial amount of information regarding preferences and technology available in the data available to us, even if we have ignored some of the pressing issues of missing data in carrying out our discussion. There are two problems with missing data in our sample. One is the gaps in the data that make it impossible to use successive observations on child quality along with input demands to estimate the production parameters directly, as in (17). We observe an indication of child quality in in 1997, along with factor demand in that year, but don’t observe an indicator of the outcome of these choices until 2002, five years later. In between these dates, input decisions have been made and levels of child quality have been determined; these input decisions depend on wage and nonlabor income draws in the intervening years, possible shocks to productivity, etc. The only way to fill in these holes is to simulate the path of all of the state variables over this period using the DGP from the model.

The other type of missing data problem faced involves nonrandom missing data on wages. As mentioned above, wages have to be generated for the years between the observed child quality levels in any event.⁹ But in addition, when one or both parents supply no time to the market, the wage offer is not observed for the period. This type of selection is particularly troublesome when preferences are treated as random in the population. In this case, seeing a parent not supply time to the market is consistent with (1) that parent’s wage offer being low, (2) the household utility function weight on that parent’s leisure being high, or (3) both (1) and (2). In order to “extrapolate” preferences and wages when a large number of households have at least one parent out of the labor force requires parametric assumptions on both processes.

The final identification issue we consider is an extremely important one that is associated with the measurement of child quality. We consider child quality k to be inherently unobservable. The law of motion of child quality involves the previous period’s level of child quality, observable inputs in the production technology, a shock (contributed through random TFP), and a vector of unknown parameters. In particular, the scale of child quality is indeterminant. Furthermore, our only direct reading on child quality comes through the test score k^* . Our conceptualization of the relationship between the observed test score and latent child quality contains inherent randomness (through the Binomial distribution) and is a function of two unknown parameters in the $p(k)$ function. So once again, a scaling issue arises in this mapping. In order to mitigate the problem of scale indeterminacy, we decided to restrict both the production technology and the mapping between k and k^* . In our estimation, we restrict TFP to be deterministic and constant (equal to 1) in every period. Thus $\sigma_R = 0$, and $\gamma_{00} = \gamma_{01} = 0$. This results in 3 fewer production function parameters being estimated. Furthermore, for the probability of answering a question on the Letter Word test correctly, we restrict $\lambda_0 = 0$ and $\lambda_1 = 1$, so that

$$p(k) = \frac{k}{1+k}.$$

These restrictions seemed to result in sensible estimates of the primitive parameters and led to what might be considered plausible degrees of dispersion in the implied k distribution. Clearly, however, other restrictions could have been imposed that would have resulted in different point estimates and inferences.

4.4 Estimator

The family data we have available consist of a sample households with observed characteristics X and includes children interviewed at various ages, where child age is indexed t .

⁹While wage and nonlabor income is gathered at every interview date, PSID interviews are conducted every 2 years at this point, and our decision periods correspond to single years. For this reason, even the wage and nonlabor income process has to be simulated between the times when the child quality measures are available.

The observed household characteristics include parental variables, such as education and ages of parents at the birth of the child. We observe for each mother and father in the household, hours worked, hours spent with children (both active and passive), and repeated measures of child quality. In addition, we observe, (accepted) wages for both parents and total household income, as described in the Data Section. Although data on some child specific expenditures are available, we do not utilize them in forming the estimator.¹⁰

The family data we utilize to estimate primitive parameters is taken from the CDS administered in 1997 and 2002 and from the standard questions concerning household income, wages, and labor supply asked of all adult household members in the PSID. For each household included in our sample we observe the following information:

<i>Variable</i>	<i>Description</i>	<i>Survey Years</i>	<i>Model Years</i>	<i>Source</i>
$\{h_{1,t}, h_{2,t}\}$	Parental labor supply	1997,1999,2001	1996,1998,2000	PSID
$\{w_{1,t}, w_{2,t}, I_t\}$	Parental wages and nonlabor income	1997,1999,2001	1996,1998,2000	PSID
$\{k_t^*\}$	Letter word score	1997,2002	1997,2002	CDS
$\{\tau_{1,t}, \tau_{2,t}, z_{1,t}, z_{2,t}\}$	Time spent with child by parent	1997,2002	1997,2002	CDS
X	Demographic characteristics	1997	1997-	PSID

The survey year column refers to the year of the survey from which the information is taken. The model year column indicates our timing convention regarding this information. The last column gives the source of the information. So, for example, the parental labor supply information is taken from the regular PSID survey given in years 1997, 1999, and 2001, and the information is retrospective (hours per week in the previous year), so that it is taken to refer to outcomes in 1996, 1998, and 2000, respectively. The demographic information, containing the (1997) age of the parents and their schooling completion levels, is taken from the 1997 PSID. Schooling is assumed to be time invariant and age is incremented in the obvious way.

The estimator is based on simulation. We first define a set of sample moments, which summarize the relationships in the sample at at each survey date and across survey dates. Let the vector of sample characteristics in our sample of size N be denoted by M_N . For each household i , we generate a set of NS sample paths over the development period in the following manner. The empirical process begins in 1997 when the child is t_i^0 years of age in household i . Given the parent’s characteristics at the sample date, X_i^0 , we draw from the distribution of shocks to wages and nonlabor incomes, and in conjunction with the “mean shifters” in X_i^0 , we determine the initial wage and income draws. Using the 1997 Letter Word test score, k^* , we draw from the distribution of k as described in Section 4.2. The particular draw of k is an initial condition for this particular sample path. Finally, we also

¹⁰The reason for this decision was that the average reports of child-specific expenditures were exceedingly low. Our interpretation was that respondents were not properly attributing some of the household expenditures on public goods to children. In any case, we think that it is difficult for any person, even an economist to properly impute these values, and hence did not utilize them when forming the estimator.

draw from the distribution of household preferences, $G(\alpha)$, and this draw stays with the household over the entire sample path. Since the production technology is assumed to be deterministic, we can then solve the household's decision problem in 1997, yielding values of labor supply and investment decisions.

In 1998 we have no observations on child quality or investments, yet using the DGP we can simulate these values. Since the wage and nonlabor income processes are assumed to be conditional (on observable characteristics) independently distributed over time, we draw new wage offers and nonlabor income for 1998. Child quality in 1998 is determined from child quality in 1997 and the inputs used in 1997. Household decisions for 1998 are then determined, and we repeat this process through the year 2002. In this manner we generate a sequence of wage and nonlabor income draws from 1997 through 2002, child quality realizations (k) in all years, and sequences of all of the other dependent variables in the model.

For the same household i , this process is repeated NS times, so that in the end we have $NS \times N$ sample paths. Using the simulated data set, we then compute the analogous simulated sample characteristics to those determined from the actual data sample. The simulated moments generated are determined by Ω , the vector of all primitive parameters that characterize the model, and the actual (pseudo) random number draws made in generating the sample paths. Denote the simulated sample characteristics generated under the parameter vector Ω by $\tilde{M}_{NS}(\Omega)$. The Method of Simulated Moments (MSM) estimator of Ω is then given by

$$\hat{\Omega}_{NS,N,W} = \arg \min_{\Omega} (M_N - \tilde{M}_{NS,N}(\Omega))' W (M_N - \tilde{M}_{NS,N}(\Omega)),$$

where W is a symmetric, positive-definite weighting matrix. By the Law of Large Numbers (LLN), $plim_{N \rightarrow \infty} M_N = M$, the population vector of values of the characteristics comprising M_N . Since we do not have access to unbiased simulators, for consistency of the MSM estimator, we require that NS also grow indefinitely large. Let the true value of the parameter vector characterizing the model be denoted by Ω_0 . Then $plim_{NS \rightarrow \infty} \tilde{M}_{NS,N}(\Omega_0) = M_N(\Omega_0)$. Given identification and these conditions,

$$plim_{N \rightarrow \infty, NS \rightarrow \infty} \tilde{\Omega}_{NS,N,W} = \Omega \text{ for all } W.$$

The moments we use include the average and standard deviation of child quality 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 care 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 compute a number of contemporaneous and lagged correlations between the observed labor supply, time with children, child quality, wages, and income. It is important to note that while we do not observe child inputs, labor supply, wages, and income in the same periods, our simulation method allows us to tractably combine moments from various points in the child development process into a single estimator.

4.5 Data

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 families, with an oversample of black and low-income families. In 1997, the PSID began collecting data on a random sample of the PSID families that have children under the age of 13 in a Child Development Supplement (CDS-I). Data were collected for up to two children per family. The CDS collects information on child development and family dynamics, including parent-child relationships, home environment, indicators of children's health, cognitive achievements, social-emotional development and time use, among other variables. The entire CDS sample size in 1997 is approximately 3,500 children residing in 2,400 households. A follow-up study with these children and families was conducted in 2002-03 (CDS-II).¹¹ These children were between the ages of 8-18 in 2003. No new children were added to the study.

Starting in 1997, children's time diaries were collected along with detailed assessments of children's cognitive development. For 2 days per week (one weekday and one weekend day), children (with the assistance of primary caregiver for very young children) 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 participation for parents: 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 construct a weekly measure of active and passive time for the mother and father by multiplying the daily hours by 5 for the weekday and 2 for the weekend day and summing the total hours for each category of time.

Children's cognitive skills are conceived broadly to include language skills, literacy and problem-solving skills and are measured with the Woodcock Johnson Achievement Test-Revised (Woodcock and Johnson, 1989). In 1997, children aged 3-5 received the Letter-Word Identification and Applied Problem sub-tests. Children aged 6 and above received Letter-Word and Passage Comprehension sub-tests as well as Applied Problems and Calculation sub-tests. In the 2002-03 (second) wave, these tests were re-administered, with the exception of the the Calculation sub-test. Given the wide range of ages to which the Letter-Word tests was administered, we use this test as our measure of child development. We use the raw scores on this exam rather than the age standardized exam. The test

¹¹An important addition in CDS-II is a set of detailed questions about childhood consumption. These questions concern the amount of money the family and others outside of the family unit paid for various types of the consumption of and the investment in the target child over the previous 12 months. These items include tuition, tutoring programs, lessons, and sports. There is also a set of questions concerning the expenditures the family or others outside the family unit make for all the children within the household. These include toys or presents, vacations, school supplies, food and clothes or shoes. There is also information on health care and childcare and related care expenses.

contains 57 items, and the range in the raw scores is 0 to 57.

For each household, we observe hours, hourly wages, and non-labor income during the 1997, 1999, 2001, and 2003 surveys reported for the previous year (1996, 1998, 2000, and 2002). The monetary values have been deflated and are all in 2001 USD. All

wage and income information is used in estimating the model, even though child investment and achievement information is only available in 1997 and 2002-03.

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 preliminary sample consists of children who (1) have valid test scores in both waves of the Child Development Supplements, and (2) are sons or daughters of the head of the household (94 percent of the original CDS sample). Because our model pertains to a single child, we drop those households who have more than one child during the sample period. In addition we drop observations with missing information on mother's or father's time with the child, those reporting (real) hourly wages above \$150 per hour in any period, and those with weekly non-labor income above \$1,000 in any period. Our total sample consists of 91 intact single child households.

Table 1 reports descriptive statistics of our selected sample. At the initial 1997 wave of the sample, the children are aged 3-12, with the average child age of 6.3 years. In 1997, the parent's average age is 37.7 for fathers and 34.5 for mothers. Average years of schooling is similar across parents at about 13 years.

Table 1 provides descriptive statistics on labor supply, wages, and income. Average hours worked by fathers is about 30 percent larger than for mothers. Average father's wages are about 27 percent larger than average mother's wages. Average non-labor income is about \$133 per week and about 64 percent of households have no nonlabor income.

Table 2 provides evidence on the allocation of parental time as the child ages. Mothers spend almost twice as much active time with the children as fathers when the child is aged 3-5. This gap in active time largely closes for older children. When the children are young, both mothers and fathers spend much more of their total child rearing time actively engaging with the child rather than in passive engagement. For older children, the parents are spending roughly equal time in passive and active engagement as the amount of active time declines for both mothers and fathers. Because of the sharp reduction in active time with the child, the average total time with the child for mothers declines substantially as the child ages. However, total time the father spends with the child increase slightly, which is due entirely to an increase in the father's passive time engagement.

Table 3 shows the pattern of mothers and fathers labor supply by child age. There is little change in the probability of working for mother's or father's as the child ages. However, average hours worked for mothers who are employed (at least 1 hour in the labor market) does increase slightly over the period while average hours worked for fathers declines slightly.

Figure 1 presents the average letter word score by the child's age. The average score

is about 5 for 3 year old children and rises to about 50 for 14 year old children. Since the scores are not age-normed, there is a substantial degree of growth, particularly at early ages. The function flattens out due to ceiling effects at the upper ages in the development period we examine.

5 Estimates

This section begins with a discussion of the estimates of the behavioral model. We then provide an assessment of within-sample model fit. We conclude the section with a number of “local” comparative statics exercises performed using the point estimates we obtained.

5.1 Household Preference Parameters

Periods are in years and the assumed planning horizon is age 16, $T = 16$. As discussed above, parents may continue to make child investments after this point but we do not explicitly model these investments and rely on our terminal period specification to capture the utility value of them. The annualized discount rate for the household is fixed at $\beta = 0.95$.

Table 4 presents the parameter estimates of the behavioral model in which it is assumed that parental preferences are time (i.e., child age) invariant, though we allow for population heterogeneity in preferences. The transformed parameters of the distribution are difficult to interpret, so instead we present their mean values, standard deviations, and the three correlation coefficients, which taken together describe the (estimated) first two moments of the population distribution of preferences under our parametric assumption on this multivariate distribution. In terms of the average household preference weights, under the second specification, there is a somewhat higher weight placed on father’s leisure than the mother’s (0.17 versus 0.13). The average weight applied to household consumption is 0.33, and the weight attached to child “quality” is 0.37. The parameter estimates are mainly attributable to the fact that women spend “too little” time in leisure given the wage differences between fathers and mothers and the declining gender differences in the productivity of time investments with the age of the child to be consistent with equal average household preference weights.

Turning next to what the estimates imply about the dispersion of preferences, the parameter estimates for Model 2 yield standard deviations for the four preference parameters of between 0.08 and 0.2. The ratio of standard deviation to mean (coefficient of variation) for the preference parameters varies between 0.28 and 0.77, with the relative dispersion greatest in the preference for mother’s leisure. We also estimate a strong correlation in leisure preferences across spouses, with the correlation in α_1 and α_2 estimated at 0.5. This may reflect the extent of assortative matching in the marriage market with regard to preferences in leisure. In contrast, we estimate a negative correlation between preferences for

mother’s leisure α_1 and household consumption α_3 , estimated at -0.25 , but a positive correlation between father’s leisure α_2 and consumption α_3 at 0.4 . These correlations are likely not representative of the general population (including households with no children and those with more than one child) but reflect the particular sample of single child households we use to estimate the parameters.

The scaling factor applied to the terminal valuation of child quality (ψ) is estimated to be approximately 43.4 . To give some meaning to this number, if we were to assume the household is infinitely lived (as discussed above) and using the assumed discount factor of $\beta = 0.95$, the implied ψ terminal value on child quality would be $(1 - \beta)^{-1} = 20$. Taken together, our point estimates imply that the household highly values child quality both in terms of flows and its terminal value, where the termination date here is considered the end of this particular (sub) development process. The high value of ψ could indicate that k_T serves as an “important” initial condition for developmental processes that begin in the later teen years. However, unlike the estimates for the other preference parameters, this parameter has a relatively high standard error of 29.13 . This suggests that it is difficult to precisely estimate the terminal value of child quality separately from the flow value, given by the α_4 .

5.2 Child Quality Technology Parameters

We next discuss the estimated production process for child quality, which we allow to change in a “smooth” way with the age of the child, as described above. Table 5 provides the parameter estimates.

Figures 2 and 3 graph the estimated technology parameters from Table 5. Figure 3 shows that the persistence of child quality (δ_{6t}) is increasing as the child reaches the upper age limit of our analysis. While we believe that this may reflect a real characteristic of the development process, there is no doubt that it also reflects the ceiling effect produced by our fixed interval measure of child quality that is not age-normed.

We see that for three of the time “flow” inputs into the dynamic production process (mother’s active and passive time and father’s active and passive time), the productivity the input changes substantially over the child development process. As expected, mother’s active time is the most productive input for young children, followed by active time of the father. For young children, passive time from mothers and fathers has much lower productivity. The productivity of mother’s and father’s active time is declining with the child’s age. While passive time for the mother stays about the same, the productivity of passive father’s time is actually increasing as the child ages. By the time the child reaches age 12, we estimate that passive father’s time is more productive than active mother’s or active father’s time. This change in the productivity of time as the child ages reflects the changing input mix of time revealed in the data. Fathers spend more time with their child as the child ages and much of this time is of a passive sort.

The declining productivity of active parental time makes some intuitive sense given our

model specification.

Once children attain the age of 5 or 6, they typically leave the home for significant periods of time each day for formal schooling activities. This amounts to a large, probably discontinuous, shift in the child quality process. Even after the onset of formal schooling activities, the child may increasingly be subject to inputs, both good and bad, from teachers and other students, that supplant the interactions the child had previously with the parents. From the point of view of parental inputs, their input decisions have increasingly small effects on child outcomes as they are “crowded out” by these others.¹² While one could argue about the form of the dependence of the production process on the age of the child, it is reasonable to think that the impact of parental inputs is, in general, declining.¹³

5.3 Wage and Non-Labor Income Process Parameters

Wage offers are unobserved with a distribution that depends on observed household characteristics, including mother’s and father’s age and mother’s and father’s education. Table 6 displays the estimated wage and income parameter estimates. We estimate that each year of schooling for mothers increases her mean log wage offer by around 4.8 percent, and for fathers, by around 7.8 percent. We estimate a rather flat but convex age earnings profile for these households, rather than a standard concave pattern. This is likely attributable to the limited age range at which we observe the parents, when parents are aged between 30 and 50. The parameter estimates imply that the mother’s mean earnings would increase by about 14 percent from age 30 to age 45, and about 18 percent for fathers. We estimate that the innovation/shock process for wage offers for mothers and fathers have substantial positive correlation at 0.8, which probably reflects both that the mother and father are in a common labor market and assortative marriage matching.

Non-labor income for the household is partially observed for some periods but unobserved for others. Unlike the wage offer distribution, we do not face the problem of endogenous non-labor income since we have an observation of non-labor income for each household. The mean of the latent non-labor income is given by μ_3^0 and does not depend on parental characteristics.¹⁴ For the latent non-labor income process, we estimate a high standard deviation in the innovation. There is some imprecision in the estimate of μ_3^0 intercept but the standard deviation of the shock is quite precisely estimated.

¹²Of course, the parents continue to have a major impact on the factor inputs through their choice of the child’s schooling environment. Liu et al. (2010) focus on this important aspect of child investment decisions.

¹³Several researchers have pointed to the importance of the phenomenon of self-investment as the child ages (e.g., Monfardini (2008)). The persistence we note in child quality process as the child ages may be due to the child, and others, supplying inputs that are unobserved and persistent.

¹⁴We found that non-labor income is in general unrelated to household characteristics of parents, education and age. Hence we model the latent distribution not to depend on these characteristics.

5.4 Within Sample Fit

Table 7 displays the sample fit of our simulated model to some features of the wage and income data. In general, we fit the mean and standard deviation of accepted wages and non-labor income well.

In terms of time allocation, the estimated model is able to fit basic patterns in time with children and labor supply. Table 8 presents the sample fit for each of the two types of parental time and the probability working and hours worked for those who work. Employment is defined as working any hours during the week of the survey (note any hours during the year). The model is able to replicate the high employment rates for fathers relative to mothers, and the higher average hours of work for fathers relative to mothers. In addition, the model reproduces the increasing average work hours for mothers as their child ages. However, while the data indicates a slight fall in probability of employment for mothers as the child ages (from 0.83 to 0.79), the model indicates an increase in employment rates.

Examining the within sample fit of time allocation to children, we see that the model fits the lower average time fathers spend with their child than mothers. In addition, the model replicates the declining time mothers spend with children as their children age, although the estimated model predicts a flatter change in active mother’s time than in the sample data.

Figure 4 provides evidence on the sample fit of the estimated child quality process to the observed child quality measure. The estimated model fits well the concave increasing average level of child quality as the child ages.

5.5 Comparative Statics Exercise: Wage Changes

Next, we probe the predictions of the model at the estimated parameters by performing a number of comparative statics exercises. A potentially interesting question is how changes in wages earned by fathers or mothers impact household decisions and outcomes. The elasticities we compute are “finite” ones. For each period t , we increase the wage draw for all mothers or all fathers. We then re-simulate the model for each sample household from the initial child interview in 1997 to the terminal period at age 16, and re-calculate simulated moments over the data sample and our simulation draws from the estimated distributions of preferences, wages, non-labor income, and child quality shocks.

Table 9 computes elasticities for mean child quality (at the terminal age of 16), mean hours worked by both parents, mean time with the children of both parents, leisure of both parents, child expenditures, parental consumption, and overall household utility. The first column provides the baseline levels at the original wage draws. The age 16 child quality is the terminal value of child quality produced by parental inputs at the last period of development, age 16, and the current level of child quality at that age.¹⁵ We focus on this

¹⁵Given our timing convention, where inputs in period t produce child quality in period $t+1$, the terminal level of child quality produced from age 16 inputs is k_{17} , with test score measure k_{17}^* .

statistic since the terminal value of child quality seems to be the most policy relevant from the perspective of producing adult outcomes. This terminal level of child quality is the initial condition for the child as she transitions to adulthood. We report both the simulated value of *measured* child quality on the test score scale of 0 to 57, and the simulated value of *latent* child quality, the value that enters the household’s utility function and the production function. The remaining statistics are summary measures over all child ages. For example, average hours that the mother works under the baseline is the simulation estimate of the average labor supply for all sample mothers who have children that range in age from 3 to 16. The elasticity estimates indicate that a 1 percent increase in the women’s wage offer would increase average labor supply for mothers by 1.05 percent from baseline (a greater than unit elastic response).

Our main interest is in the difference in the responsiveness of each of the variables to wage changes by gender. Since our main concern is child quality, we begin with the first two rows of Table 9. We note that increasing the mother’s wage has a positive impact on child quality. The elasticity of latent child quality to mother’s wage is 0.0049, a very inelastic response. Given the ceiling effect in the test score measure, the elasticity of the measure child quality is even lower. The net effect of this change depends on preferences over consumption, parental leisure, and child quality, and various time and expenditures choices. The remaining rows of the simulation table unpacks the various household responses. As we see from the fourth row, the mother reduces her active and passive time with the child as a result of a wage increase. This effect alone would reduce child quality. However, the higher wage for mothers induces higher labor supply and increases labor income for the household. With this added labor income, the household optimally increases its child expenditures. This effect increases child quality. In addition, the increase in the mother’s wage offer changes the relative prices of mother’s and father’s time, and the father works less and spends more time with his child. This effect of greater father’s time with the child has a positive effect on child quality. Together these positive and negative effects of an increase in the mother’s wage largely cancel each other out, resulting in a weak inelastic effect on child quality.

In the next policy experiment (column 3 of Table 9), we increase the father’s wage. An increase in the wage of the father has a stronger effect on final child quality than an increase in the mother’s wage, but the father’s child quality elasticity is still quite inelastic at 0.0067. Notice that the wage increase for fathers has a smaller negative own effect on father’s labor supply and a stronger positive cross-effect on mother’s labor supply through an income effect. An increase in the father’s wage results in the mother spending less time in the labor market, and this reduction is spread between her increased leisure and increased time with her child, which results in an improvement in child quality. The increase in the father’s wage also results in a larger increase in labor income and higher child expenditures than an increase in the mother’s wage. However, it should be noted however that given that fathers have higher average wages, a given increase in the father’s wage offer is a larger wage effect than is that for mother’s. This is reflected in the higher increase in household

utility from the increase in father’s wage relative to the mother’s wage.

5.6 Input Allocations

We investigate the importance of modeling the child development process within a household framework by considering a number of special cases of our more general model.

5.6.1 Child Quality Maximizing Preferences

In the first special case, we examine the optimal level of child inputs if we were to assume the household only has preferences over child quality. The optimal allocation of inputs under these “child quality maximizing” preferences is given by solving the dynamic household problem setting the weight on parental leisure and consumption to zero: $\alpha_1 = \alpha_2 = \alpha_3 = 0$. This is equivalent to assuming the parents solve the following maximization process:

$$\max_{\tau_{1t}, \tau_{2t}} f(k_t, \tau_{1t}, \tau_{2t}, e_t),$$

subject to the constraints: $TT = h_{jt} + \tau_{jt}$ for $j = 1, 2$ and $e_t = w_{1t}h_{1t} + w_{2t}h_{2t} + I_t$. Note that the under these “child quality maximizing” preferences, the household problem is strictly a static problem as the optimal allocation of inputs to maximize child quality in each period also maximizes the terminal period T level of child quality.

Table 10 presents the mean level of endogenous choices and final period (age 16) child quality under the baseline unrestricted model, with heterogeneous preferences discussed above, and the restricted child quality maximizing preferences. The first two rows indicate that measured and latent child quality increases under the child quality maximizing preferences, with a greater response in latent child quality. With child quality maximizing preferences, average latent child quality increases by about 1/3, from 11.15 to 15.15. With these preferences, the parents choose to work in the labor market substantially fewer hours than under the baseline, and the time spent with their children increases substantially. It is interesting to note that mothers and fathers still work under the child quality maximizing preferences and do not spend their entire time endowment on child rearing. This is because market work funds child goods expenditures, and it is optimal for the parents to continue working some hours. Even under the child quality maximizing preferences, the relative specialization of mothers to child rearing and fathers to market work still occurs as mothers continue to spend more time on average with their children than fathers, and fathers work more than mothers. Even with the lower labor income under the child maximizing preferences, expenditures on children still increases almost three fold over the baseline. By definition all of the household’s income is spent on child goods.

5.6.2 Selfish Parents Preferences

The next special case we consider is the opposite of the previous one. Here we set household preferences so that there is no weight on child quality: $\alpha_4 = 0$. With these “selfish parents preferences,” the parents solve the following problem:

$$\max_{h_{1t}, h_{2t}} U(l_{1t}, l_{2t}, c_t),$$

subject to the constraints: $TT = h_{jt} + l_{1t}$ for $j = 1, 2$ and $c_t = w_{1t}h_{1t} + w_{2t}h_{2t} + I_t$. We maintain the same distribution of preferences over leisure and consumption as estimated. These selfish parents preferences imply that parents spend no time with their children $\tau_{1t} = \tau_{2t} = 0$ for all t , and there are no expenditures on child goods $e_t = 0$ for all t . As with the child quality maximizing preferences, the household problem under the selfish parent preferences is a strictly static problem. These preferences are essentially the preferences that would prevail for couples without children, although the preference distribution over leisure and consumption we estimate is for a sample of families that choose to have a child and may not be representative of single, no child households.

The third column of Table 10 presents the results for the model estimated under these selfish parent preferences. By definition, child quality, time with children, and expenditures are zero. The interesting aspect of this special case is the contrast of the labor supply decision with that for the baseline case. If the parents place no weight on child quality, the mother and father would both work substantially more hours than under the baseline where they derive utility from child quality. As a result of the additional labor income and no child expenditures, household consumption increases. These results suggest that parents optimally reduce their labor supply in order to invest in child quality through time spent with their children. With no weight on child quality, parents would spend considerably more time in the labor market in order to finance their own consumption. The most substantial change is for women, as their average labor supply when the household puts no weight on child quality is much closer to that of their husbands (52 vs. 58) than under the baseline preferences (33 vs 43). Our estimates are in line with the view that a major reason for gender specialization of time between market work and childcare is due to household preferences over child quality. Without a preference for children, men’s and women’s labor supply would be substantially more similar.

5.6.3 Technology Optimal Allocations

In the third special case we consider, the optimal level of child inputs is determined by the technology alone. This special case is intended to replicate one of the main approaches to child investment taken in the current literature, in which the production technology alone is used to draw inferences about the optimal allocation of child inputs across stages in the child’s development and across different types of inputs. Our estimate of the production technology provides a “selection corrected” estimate of the production technology because

we jointly estimate the production technology with our model of the endogenous parental inputs. We can then use the estimated technology to indicate what the optimal allocation of inputs would be if a social planner were to use this technology to allocate resources but ignore household preferences and differential wage offers. The degree to which this technologically-optimal allocation would differ from the actual allocation parents make demonstrates the importance of considering the child quality technology jointly within the household decision context, including preferences, wages, and the non-labor income processes.

With no resource constraint, we cannot consider the level of inputs using the estimated production technology alone. Instead we define the “technology optimal” ratio of inputs as the ratio of marginal productivities of the inputs in each period. The technology optimal ratio of mother’s time to father’s time is

$$\frac{\tau_{1t}}{\tau_{2t}} = \frac{\delta_{1t}}{\delta_{2t}},$$

and the technology optimal ratio mother’s time to child good expenditures is given by

$$\frac{\tau_{1t}}{e_t} = \frac{\delta_{1t}}{\delta_{5t}}$$

The difference between the optimal ratio of inputs chosen by the household in our unrestricted baseline model and the technology optimal ratio is that the household optimal allocation takes into account the cost of child investments from foregone parental leisure and consumption (given by $\alpha_1, \alpha_2, \alpha_3$) and the differential opportunity cost of mother’s and father’s time due to different wage offers.

Figure 5 displays the optimal ratio of mother’s time to father’s time with the child under i) the baseline unrestricted household-optimal model and ii) the technology-optimal model. For the household optimal model, we plot the mean input choices in our sample at each child age.

Figure 5 shows that the technology optimal allocation reflects the change in productivity of mothers and fathers as the child ages (Figure 2). The technology optimal allocation allocates relatively *less* mother’s time to child development than the household optimal solution. The household optimal allocation takes into account that mothers are less productive in the labor market than fathers, and therefore the optimal allocation of time is for mothers to spend relatively more time with the child than fathers. Both allocations are downward sloping given that the main time invariant feature of the model is the technology. Mother’s time is becoming increasingly less productive relative to father’s time as the child ages.

Figure 6 repeats the analysis of optimal allocations focusing on another ratio of inputs, the ratio of mother’s time to child good expenditures. The technology optimal allocation of mother’s time to expenditures is several times higher than the baseline household optimal ratio. This reflects the fact that the household optimally allocates a much higher relative

level of expenditures to the child than would be indicated solely by technological considerations. The household optimally substitutes child goods for mother's time. The technology optimal solution ignores the fact that mother's time with the child has both an opportunity cost in terms of foregone leisure for the mother and foregone labor income from mother's labor supply, and therefore foregone parental consumption and child expenditures. This difference is particularly stark when children are young, as the technology alone would dictate that mothers spend substantially more time raising children than the household would optimally choose. This difference declines as the child ages and the productivity of mother's time falls.

6 Social Policy and Child Development

In this section we consider, in a highly stylized manner, the impact of two social policies on child development and household welfare. Our entire analysis has been conducted outside of an equilibrium framework, so the exercises reported here are not to be considered part of a serious policy analysis. However, given that one could reasonably argue that general equilibrium effects would tend to attenuate the impacts found here, our exercise might at least establish some bounds on the potential impacts of these policies on child development.

6.1 Public Schooling

Given that the developmental period we examine covers the age range 3-16, school attendance occurs over most of the period. We have not accounted for this in the model explicitly, which may arguably have resulted in model misspecification regarding the choice sets of parents and inputs in the child development process.¹⁶ In the United States, most households with children 5 years of age or less receive minimal (child-related) contributions from governmental and nongovernmental organizations. At the age of 5 or 6, virtually all children enter public schooling; within this institutional environment, children and their parents receive large transfers from state and local governments, primarily in the form of teachers, administrators, school buildings, books, subsidized meals, etc. Data from the National Center for Educational Statistics suggests that the average of state expenditures per pupil (across all grade levels) was approximately 10,000 dollars in 2005. We use this amount in our exercise, and assume that this translates into a weekly transfer of \$192.¹⁷

Of course, there are many potential benefits of public schooling aside from indirect transfers to households with school-age children. The most obvious ones are efficiency gains, since there would seem to exist scale economies in the provision of basic education,

¹⁶Expanding the set of inputs to include time spent with teachers and others encountered in institutional settings, along with the monetary contributions made by such actors, is an objective of our current research.

¹⁷We have simply expressed the total school year amount as a weekly flow value. Our model with its conventions regarding decision periods (the year) and measurement units (weeks) does not allow us to allocate the transfers only to weeks in which school is in session.

and some studies have pointed to the positive impacts of group interactions when young on noncognitive skills later in life (while there is a healthy debate concerning the deleterious effects of classroom overcrowding, few researchers or parents call for class sizes of less than 10 students). We ignore all of these important arguments for the benefits of public schooling, and here examine only the implications of these indirect transfers for the welfare of children and households using our modeling framework.

In terms of the household's problem, in our model public schooling simply amounts to a contribution to e , expenditures on the child. The difference between this form of subsidy and a direct transfer of \$192 to the household with no restrictions on how it is to be spent are obvious. Household welfare can never be increased by restricting the manner in which the transfer is used, whereas child quality could be. Given our functional form assumptions regarding child quality production, every household spends some money on child investment goods, e . If without the public schooling subsidy, the household spent \$10 per week on the child, say, then the subsidy takes the household well beyond this level. In terms of household welfare, the constraint that at least \$192 of total household income be spent on child quality would be a binding one. The household will be better off than it would be without the transfer, but not as well off as if the transfer were unrestricted. Child quality could increase appreciably in such a case, however. In contrast, when the household was already spending more than \$192 before receiving the public school subsidy, the constraint is non-binding. The household will be better off, and a portion of its higher income will be spent on child quality investment. In this case, the percentage gain in child quality may not be as large as when the constraint is binding.

Table 11 displays the results from this policy experiment. We computed the results using the same technique as described above for the wage offer comparative statistics exercises. Column 2 reports the results from a transfer of income of \$192 to each household for all periods subject to the constraint that at least \$192 must be spent on child investment goods. We see that this policy increases average terminal (latent) child quality by 3.6 percent over the baseline. Given the ceiling on the test score measure of 57, measured child quality increases by less than 1 percent. Mothers and fathers react to this policy by reducing their labor supply. This occurs because of the income effect of the increase in household income. On average, mothers and fathers increase their time spent with their children, but the average increase is less than the average reduction in labor supply. We can see then that the reduction in labor supply caused by the income transfer is being split between greater leisure for the parents and greater time investment in their children. Child good expenditures increase quite substantially, by almost 30 percent, as a result of the targeted transfer. Household consumption also increases because of the income transfer, but far less than the increase in child good expenditures. For unconstrained households, those households already spending in excess of the required \$192 on child goods, this policy is a pure income transfer.

6.2 Limitations on Work Hours

In France in 2000, the left-wing government of Lionel Jospin passed a law establishing 35 hours as the new legal work week, with strict limits on and penalties for work over this limit.¹⁸ While the rationale for decreasing the length of the work week was not framed explicitly in terms of child development, such a law could possibly be expected to have significant impacts on the child development process, since it directly impacts the household's ability to freely allocate time and financial resources across competing uses. Needless to say, if such a law were to be imposed and seriously enforced in the U.S. context, we would expect to see potentially major changes in wage distributions, the demand for capital, etc. These are all neglected in our exercise, where we simply impose a binding 35 hour work limitation on each of the parents and make no other changes to the choice problem the household faces. We assume that the law is perfectly enforced and permanent.

Table 11 provides the results from this policy experiment in the last column. Under this policy, average latent and measured child quality increases by less than 1 percent. Mean hours of work for the mother and father fall precipitously, though by far the greatest change is for fathers, whose average labor supply falls 22.6 percent. Due to these hours restrictions, both parents spend more substantially more active and passive time with their child, on average. Mirroring the labor supply changes, the most substantial increase is for fathers. Child good expenditures fall 12.8 percent from the baseline because of the lower labor supply and consequent reduction in labor income. The positive effects of the increase of parental time spent with the child is almost perfectly countered by the negative impact of the reduced expenditure on the child. We see from the last row in the table that household utility drops by a bit more than 3 percent due to the hours restrictions. With a small positive impact on average child quality, this policy does not seem to have much to recommend it, at least evaluated in terms of our simple partial-partial equilibrium model.

7 Conclusion

Using a simple dynamic production technology for child quality and a Cobb-Douglas specification of a household utility function, we employ unique data from the PSID-CDS on investments in children to recover estimates of the parameters necessary to characterize the child development process. We view the main message of the paper to be that household time and money investments in children can only be properly understood when preferences, technologies, and choice sets are simultaneously considered and estimated. We found that while the average household attaches a substantial weight to child quality in its utility function, it is by no means only concerned with the quality of the child. This was made abundantly clear in the comparative statics exercises and policy experiments. For example, we provided quantitative evidence regarding the extent of "crowding out" of household

¹⁸The limitations on overtime hours were substantially relaxed by amendments introduced over the period 2004 through 2007.

investment in children by public investment, in the form of schooling transfers. Though we have not allowed, at present, for time inputs of teachers and other agents, it is clear that these factors of production belong in the production function as well, and we can expect that the time inputs of parents to be supplanted by those of other agents when the child begins formal schooling. Policy-makers considering changes in the school environment should be aware of parental responses to changes in schooling inputs or their attempts at enhancing child welfare are likely to be frustrated.

We were also able to reanalyze the issue of the impact of mother’s employment on their children’s development. In the comparative statics exercise in which we increased mothers’ wage offers over the development process, we found essentially no effect on child quality at the end of the developmental stage we analyze. We were able to trace out the complicated household substitution patterns that led to this result. Estimates of the child production technology and household preferences were necessary to carry out this exercise and to arrive at this conclusion.

We believe that we have uncovered a number of interesting empirical results and interpretations, though obviously much more needs to be done on both the preferences and technology sides of the model. We mentioned that the set of inputs should be expanded to include the time and money contributions of teachers, tutors, care-givers, and other relatives. Our analysis is also limited in the sense that we specify a unitary model to represent household preferences. We used a household utility function approach to ease the computational burden and to focus attention on the child quality production function. We do not consider the specification of a unitary model to be that problematic in the present context, given the difficulty of convincingly estimating individual preferences with the data available to us. However, it is clear that to include in our analysis single parent households, a decidedly important objective, it would be necessary to abandon the household utility function framework for an individual-based objective alternative, and to modify characteristics of the production technology we have employed here as well.¹⁹

Despite the limitations of the analysis, we believe that we have been able to demonstrate the importance of a variety of parental investment activities in the child development process, and that we have demonstrated the extent to which these investment activities respond to differences in resource constraints and preferences both across households and over time. While many simplifying assumptions were made in order to ease the computational burden and make identification more transparent, the model was able to fit the data in a satisfactory manner. We view the results here as a potentially important first step in building an equilibrium model of the child development process that can be used in designing social and family policies to increase average “child quality” while decreasing its dispersion in the population.

¹⁹Our production technology carries the implication that unless both parents supply time to child investment, the child quality will be zero in all subsequent periods. We would have to modify the production technology accordingly if we considered households in which one parent supplied no time to the investment process.

References

- [1] Almond, Douglas., and Janet Currie 2010.”Human Capital Development Before Age Five,” NBER Working Papers 15827
- [2] Averett, Susan, Lisa Gennetian, and Elizabeth Peters. 2005. “Paternal Child Care and Children Development.” *Journal of Population Economics*, 18: 391-414.
- [3] Baydar, Nazli and Jeanne. Brooks-Gunn. 1991 “Effects of Maternal Employment and Child Care Arrangements on Preschoolers’ Cognitive and Behavioral Outcomes: Evidence from Children of the National Longitudinal Survey of Youth.” *Developmental Psychology*, 27: 932–45
- [4] Becker, Gary. 1981. *A Treatise on the Family*. Cambridge, MA: Harvard University Press.
- [5] Becker, Gary. S., and N. Tomes (1979): An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility,” *Journal of Political Economy*, 87: 1153-89
- [6] Bernal, Raquel. 2008. “The Effect Of Maternal Employment And Child Care On Children’S Cognitive Development.” *International Economic Review*, 49(4): 1173-1209.
- [7] Bernal, Raquel, and Michael P. Keane. 2010. “Quasi Structural estimation of a model of child care Choices and child cognitive ability production“ *Journal of Econometrics* vol 156 (1):164-189.
- [8] Blau Francine D., and Adam Grossberg. 1992. “Maternal Labor Supply and Children’s Cognitive Development.” *Review of Economics and Statistics*, 74: 474-481.
- [9] Blau, David M. (1999). ”The Effect of Income on Child Development,” *The Review of Economics and Statistics*,81:261-276.
- [10] Brooks-Gunn, Jeanne, Wen-Jui Han, and Jane Waldfogel. 2002. “Maternal Employment and Child Cognitive Outcomes in the First Three Years of Life: The NICHD Study of Early Child Care.” *Child Development*, 73(4): 1052-1073.
- [11] Cunha, Flavio and James J. Heckman. 2008. “Formulating and Estimating the Technology of Cognitive and Noncognitive Skill Formation.” *Journal of Human Resources*,. 43: 738-778.
- [12] Cunha, Flavio and James J. Heckman Susan Schennan.2010. “Estimating the technology of Cognitive and Non Cognitive Skill Formation“ *Econometrica* vol 78(3):883-931

- [13] Dahl, Gordon., and Lance Lochner 2008. "The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit," NBER Working Papers 14599.
- [14] Desai, Sonalde, Lindsay P. Chase-Lansdale and Robert T. Michael 1989. "Mother or Market? Effects of Maternal Employment on four-year-olds' intellectual abilities" *Demography* 26:545-561
- [15] Dooley, Martin, and Jennifer Stewart 2004. "Family Income and Child Outcomes in Canada," *Canadian Journal of Economics*, 37, 898-917.
- [16] Ermisch, John and Marco Francesconi. 2005. "Parental Employment and Children's Welfare," in *Women at Work: An Economic Perspective* (Boeri, Tito, Daniela Del Boca, and Christopher Pissarides). Oxford: Oxford University Press.
- [17] Hill, Martha. S., Yeung, Wei Jeanne., Duncan, Greg.J. 2001. "The Effect of Family Structure on Children's Economic and Behavioral Outcomes." *Journal of Population Economics*, 14: 271-299.
- [18] Hofferth, Sandra.L., Davis-Kean, P., Davis, J., Finkelstein, J. 1998.. *User guide for the Child Development Supplement of the Panel Study of Income Dynamics*. <http://psidonline.isr.umich.edu/CDS/usergd.html>.
- [19] Haveman, Robert, and Barbara Wolfe. 1995. "The Determinants of Children's Attainments: A Review of Methods and Findings." *Journal of Economic Literature*, 33: 1829-1878.
- [20] Huston, Aletha and Stacey Aronson. 2005. "Mothers' Time with Infant and Time in Employment as Predictors of Mother-Child Relationships and Children's Early Development." *Child Development*, 76: 467-482.
- [21] Hsin, Amy 2008 "Mothers' Time with Children and the Social Reproduction of Cognitive Skill" mimeo.
- [22] Liu, Haiyong, Thomas Mroz, and Wilbert Van der Klaauw. 2010. "Maternal Employment, Migration, and Child Development." *Journal of Econometrics*, 156(1): 212-228.
- [23] Levy, Daniel M., and Greg Duncan 2000 "Using Sibling Samples to Assess the Effect of Childhood Family Income on Completed Schooling," JCPR Working Papers 168, Northwestern University.
- [24] Loken, Katrine 2010. "Family Income and Children's Education: Using the Norwegian Oil Boom as a Natural Experiment," *Labour Economics*, 17, 118-29
- [25] Loken, Katrine, Magne Mogstad, Matt Wiswall 2010. "What Linear Estimators Miss: Re-Examining the Effects of Family Income on Child Outcomes" working paper.

- [26] Mayer, Susan. 1997. "What Money Can't Buy: Family Income and Children's Life Chances". Harvard University Press.
- [27] Lundberg, Shelly. 2005. "Sons, Daughters, and Parental Behavior." *Oxford Review of Economic Policy*, 21(3): 340-356.
- [28] Waldfogel, Jane, Han, Wen-Jui, and Jeanne Brooks-Gunn. 2001. "The Effects of Early Maternal Employment on Child Development." *Demography*, 39: 369-392.
- [29] Todd, Petra and Ken Wolpin. 2003. "On the Specification and Estimation of the Production Function for Cognitive Achievement." *Economic Journal*, 113: F3-F33.
- [30] Ruhm, Christopher J. 2004. "Parental Employment and Child Cognitive Development." *The Journal of Human Resources*, 39 (1): 155-192.
- [31] Vandell, Deborah and Jean Ramanan. 1992. "Effects of Early and Recent Maternal Employment on Children from Low-income Families." *Child Development*, 63: 938-49.
- [32] Yeung, Jean, Greg Duncan, and Martha Hill. 1999. "Putting Fathers Back in the Picture." *Journal of Marriage and the Family*, 49(2): 97-114.

Table 1: **Descriptive Statistics**

1997 PSID-CDS		
	Mean	Std.
Child's age	6.31	2.93
Fraction Male	0.539	0.499
Mother's age	34.50	6.48
Father's age	37.67	8.48
Mother's education	13.63	2.22
Father's education	13.62	2.06
Mean Letter Word raw score	23.89	16.21
Median LW raw score	20	
Minimum LW raw score	3	
Maximum LW raw score	55	
1996-2002 PSID		
	Mean	Std.
Mother's Work Hours per Week	31.31	16.30
Father's Work Hours per Week	44.17	10.45
Mother's Hourly Wage	15.70	10.80
Father's Hourly Wage	19.89	12.26
Non-Labor Income per Week	133.10	213.96

Notes: Sample of intact households (mother and father present in household) with one child. The top panel statistics are for the year 1997 from the 1997 PSID-CDS. Work hours, wages, and non-labor income statistics are averaged over all years of PSID data.

Source: PSID-CDS combined sample from 1997 and 2002 interviews and 1997, 1999, 2001, 2003 PSID core data.

Table 2: **Parent's Time with Child by Child Age**

Child Age	Mother	Father
	Active Time (Avg.)	
3-5	32.65	17.91
6-8	20.86	15.59
9-11	19.45	12.59
12-15	16.12	12.55
	Passive Time (Avg.)	
3-5	16.40	7.90
6-8	12.65	10.25
9-11	13.83	10.45
12-15	17.99	15.25

Notes: Sample of intact households (mother and father present in household) with one child.

Source: PSID-CDS combined sample from 1997 and 2002 interviews and 1997, 1999, 2001, 2003 PSID core data.


Table 3: **Parent's Labor Supply by Child Age**

Child Age	Mother	Father
	Fraction Working > 0 Hours	
3-5	0.831	0.985
6-8	0.829	0.991
9-11	0.862	0.977
12-15	0.794	0.985
	Average Hours Working	
3-5	30.67	45.13
6-8	35.91	44.98
9-11	36.63	45.12
12-15	34.67	43.89

Notes: Sample of intact households (mother and father present in household) with one child.

Source: PSID-CDS combined sample from 1997 and 2002 interviews and 1997, 1999, 2001, 2003 PSID core data.

Figure 1: Average Letter Word Score by Child Age
Mean Child Quality by Age



Notes: Sample of intact households (mother and father present in household) with one child.

Source: PSID-CDS combined sample from 1997 and 2002 interviews and 1997, 1999, 2001, 2003 PSID core data.


Table 4: **Preference Parameter Estimates**

	Estimate	SE
Mean of α_1	0.130	(0.0067)
Mean of α_2	0.169	(0.0050)
Mean of α_3	0.329	(0.0101)
Mean of α_4	0.373	(0.0054)
Std. of α_1	0.104	(0.0105)
Std. of α_2	0.081	(0.0044)
Std. of α_3	0.092	(0.0074)
Std. of α_4	0.205	(0.0101)
Correlation of α_1 and α_2	0.492	(0.109)
Correlation of α_1 and α_3	-0.254	(0.089)
Correlation of α_2 and α_3	0.403	(0.088)
ψ (Terminal Payoff to Child Quality)	43.67	(29.13)

Table 5: **Technology Parameter Estimates**


	Estimate	SE
γ_{10} (Mother's Active Time intercept)	-1.52	(0.028)
γ_{11} (Mother's Active Time slope)	-0.082	(0.0042)
γ_{20} (Father's Active Time intercept)	-1.81	(0.055)
γ_{21} (Father's Active Time slope)	-0.046	(0.0043)
γ_{30} (Mother's Passive Time intercept)	-2.48	(0.035)
γ_{31} (Mother's Passive Time slope)	0.0050	(0.0043)
γ_{40} (Father's Passive Time intercept)	-2.69	(0.041)
γ_{41} (Father's Passive Time slope)	0.042	(0.0052)
γ_{50} (Child Expenditures intercept)	-3.41	(0.022)
γ_{51} (Child Expenditures slope)	0.085	(0.0076)
γ_{60} (Last Period's Child Quality intercept)	-2.079	(0.065)
γ_{61} (Last Period's Child Quality slope)	0.040	(0.0056)

Figure 2: Estimated Child Development Parameters by Child Age



Notes: $\delta_{1t}, \delta_{2t}, \delta_{3t}, \delta_{4t}$ estimated parameters by child age (from Table 5).

Figure 3: Estimated Child Development Parameters by Child Age



Notes: δ_{5t} and δ_{6t} estimated parameters by child age (from 5).

Table 6: **Wage and Income Parameter Estimates**

	Estimate	SE
Mother's Log Wage Offer		
μ_1^0 (Intercept)	1.58	(0.017)
μ_1^1 (Mother's Education)	0.048	(0.0013)
μ_1^2 (Mother's Age)	0.0043	(0.0004)
μ_1^3 (Mother's Age Sq x 1000)	0.12	(0.018)
σ_1 (Standard Deviation of Innovation)	0.044	(0.032)
ρ_{12} (Correlation with Father's Wage Shock)	0.729	(0.018)
Father's Log Wage Offer		
μ_2^0 (Intercept)	1.37	(0.023)
μ_2^1 (Father's Education)	0.078	(0.0016)
μ_2^2 (Father's Age)	0.0082	(0.0006)
μ_2^3 (Father's Age Sq x 1000)	0.042	(0.017)
σ_2 (Standard Deviation of Innovation)	0.755	(0.058)
Latent Non-Labor Income		
μ_3^0 (Intercept)	-50.81	(35.31)
σ_3 (Standard Deviation of Innovation)	397.74	(33.07)

Table 7: **Sample Fit for Wages and Income**

	Data	Simulated
Avg. Mother's Wage	15.70	15.89
Std. Mother's Wage	10.82	10.08
Avg. Father's Wage	19.89	19.75
Std. Father's Wage	12.28	12.14
Avg. Non-Labor Income	133.10	133.10
Std. Non-Labor Income	213.96	213.96
Fraction with 0 Non-Labor Income	0.637	0.651

Notes: Data is actual data from sample of intact households (mother and father present in household) with one child. Simulated is the model prediction at estimated parameters given above.

Source: PSID-CDS combined sample from 1997 and 2002 interviews and 1997, 1999, 2001, 2003 PSID core data.


Table 8: **Sample Fit of Mother and Father’s Time Allocation by Child Age**

Probability Work > 0 Hours					
		Mother		Father	
Child Age	Data	Simulated	Data	Simulated	
3-5	0.831	0.829	0.985	0.996	
6-8	0.829	0.851	0.991	0.9995	
9-11	0.862	0.866	0.977	0.998	
12-15	0.794	0.880	0.985	0.997	
Hours Worked if Work (Avg.)					
		Mother		Father	
Child Age	Data	Simulated	Data	Simulated	
3-5	30.67	30.69	51.53	45.13	
6-8	35.91	30.52	50.81	44.98	
9-11	36.63	30.93	42.67	45.12	
12-15	34.67	35.60	48.20	43.89	
Active Time (Avg.)					
		Mother		Father	
Child Age	Data	Simulated	Data	Simulated	
3-5	32.65	26.39	17.91	17.04	
6-8	20.86	18.76	15.59	14.42	
9-11	19.45	17.15	12.59	13.40	
12-15	16.12	17.82	12.55	16.46	
Passive Time (Avg.)					
		Mother		Father	
Child Age	Data	Simulated	Data	Simulated	
3-5	16.40	12.80	7.90	8.94	
6-8	12.65	12.85	10.25	10.72	
9-11	13.83	14.04	10.45	11.93	
12-15	17.99	20.57	15.25	20.70	

Notes: Data is actual data from sample of intact households (mother and father present in household) with one child. Simulated is the model prediction at estimated parameters given above.

Source: PSID-CDS combined sample from 1997 and 2002 interviews and 1997, 1999, 2001, 2003 PSID core data.

Figure 4: Simulated and Actual Average Letter Word Score by Child Age



Notes: Data is actual data from sample of intact households (mother and father present in household) with one child. Simulated is the model prediction at estimated parameters given above.

Source: PSID-CDS combined sample from 1997 and 2002 interviews and 1997, 1999, 2001, 2003 PSID core data.

Table 9: **Comparative Statics**

	Level at	Increase	Increase
		in Mother's Wage	in Father's Wage
	Baseline	Elasticity	
Mean Measured Child Quality (Age 16)	51.99	0.0005	0.0007
Mean Latent Child Quality (Age 16)	11.15	0.0052	0.0068
Mean Hours Work (Mother)	33.02	1.052	-1.023
Mean Hours Work (Father)	43.32	-0.59	0.59
Mean Active Time w/ Child (Mother)	19.39	-0.49	0.49
Mean Active Time w/ Child (Father)	15.76	0.39	-0.39
Mean Passive Time w/ Child (Mother)	16.76	-0.49	0.49
Mean Passive Time w/ Child (Father)	15.43	0.39	-0.39
Mean Leisure (Mother)	42.84	-0.40	0.38
Mean Leisure (Father)	37.49	0.36	-0.36
Mean Child Expenditures	241.25	0.40	0.57
Mean Household Consumption	1,325	0.39	0.59
Mean Utility	–	0.13	0.19


Notes: This table reports elasticity estimates from an increase in mother's (2) and father's wage offer (3). Mean Measured Child Quality (Age 16) is the average level of measured child quality (test score on scale of 0 to 57) at the end of age 16 or the start of period $t = 17$, k_{17}^* . 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 10: **Optimal Decisions with Alternative Preferences**

	Level at Baseline	Child Quality Maximizing Preferences	Selfish Parent Preferences
Mean Measured Child Quality (Age 16)	51.99	53.41	0
Mean Latent Child Quality (Age 16)	11.15	15.15	0
Mean Hours Work (Mother)	33.02	8.33	51.82
Mean Hours Work (Father)	43.32	17.73	57.53
Mean Active Time w/ Child (Mother)	19.39	56.86	0
Mean Active Time w/ Child (Father)	15.76	49.56	0
Mean Passive Time w/ Child (Mother)	16.76	46.81	0
Mean Passive Time w/ Child (Father)	15.43	44.71	0
Mean Leisure (Mother)	42.84	0	60.18
Mean Leisure (Father)	37.49	0	54.47
Mean Child Expenditures	241.25	672.13	0
Mean Household Consumption	1,325	0	2,139


Notes: Child Quality Maximizing Preferences set the preference weight on parental leisure and consumption to 0: $\alpha_1 = \alpha_2 = \alpha_3 = 0$. Under these preferences, the household then maximizes the level child quality, and consumption $c_t = 0$ for all t . Selfish Parent Preferences set $\alpha_4 = 0$, and the household puts no weight on child quality. With these preferences, all child inputs equal 0 for all t . Mean Measured Child Quality (Age 16) is the average level of measured child quality (test score on scale of 0 to 57) at the end of age 16 or the start of period $t = 17$, k_{17}^* . 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} .

Figure 5: Optimal Ratio of Mother's and Father's Active Time with Child under Different Modeling Assumptions



Notes: Technology Optimal is the optimal ratio of active time with mother versus active time with father given by the technology alone: δ_{1t}/δ_{2t} . Household Optimal is the average input choices of the household under the full unrestricted model: τ_{1t}/τ_{2t} .

Figure 6: Optimal Ratio of Mother's Time with Child and Child Expenditures under Different Modeling Assumptions



Notes: Technology Optimal is the optimal ratio of active time with mother versus household expenditures on child goods given by the technology alone: δ_{1t}/δ_{5t} . Household Optimal is the average input choices of the household under the full unrestricted model: τ_{1t}/e_t .

Table 11: **Policy Simulations**

	Experiment	(1)	(2)
	Level at Baseline	Child Expend. Floor $e \geq \$192$ & Inc. Transf. of \$192	Hours Ceiling $h_{1t} \leq 35$ and $h_{2t} \leq 35$
		Percentage Change from Baseline	
Mean Measured Child Quality (Age 16)	51.99	0.36	0.08
Mean Latent Child Quality (Age 16)	11.15	3.64	0.77
Mean Hours Work (Mother)	33.01	-7.40	-15.97
Mean Hours Work (Father)	43.34	-6.24	-22.62
Mean Active Time (Mother)	19.39	4.02	8.92
Mean Active Time (Father)	15.76	4.73	13.45
Mean Passive Time (Mother)	16.76	4.30	9.02
Mean Passive Time (Father)	15.43	4.97	11.76
Mean Leisure (Mother)	42.84	2.20	4.74
Mean Leisure (Father)	37.48	3.18	15.66
Mean Child Expenditures	241.26	29.57	-12.77
Mean Household Consumption	1,326	3.17	-20.35
Mean Utility	—	3.54	-3.21

Notes: Experiment 1 gives each household an income transfer of \$192 each week and forces the household to spend at least \$192 on child goods. Experiment 2 imposes a labor supply ceiling on the parents and forces mothers and fathers to work 35 or fewer hours each week. Mean Measured Child Quality (Age 16) is the average level of measured child quality (test score on scale of 0 to 57) at the end of age 16 or the start of period $t = 17$, k_{17}^* . 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} .