The Effect of Parental Composition on Investments in Children when Markets are Incomplete.

Brant Abbott*

Yale University

This Draft: October 20, 2014

Abstract

In this paper I estimate a model of parenting decisions and child development that captures the disadvantages faced by single-parent households. In addition to having lower average resources, single-parent families are also affected by the imperfect substitutability of human capital development inputs and their lack of spousal labor supply as a source of insurance against economic risk. In the model, those who grow up with a single-parent have wages that are 37.4% lower on average, which is very similar to the data. After accounting for parental background and differences in resources the gap in the model is only 15.3%. The GMM estimates of the production function for investments in children reveal that parental time and market purchased goods and services are imperfect substitutes, and, to a lesser extent, that mother's and father's time inputs are imperfect substitutes. Imperfect substitutability of market purchased inputs and parental time interacts with uninsured economic risk, and this accounts for nearly half of the residual gap in outcomes that exists after accounting for differences in resources. Imperfect substitution of mother's and father's time only accounts for 5% of the residual gap, with the remainder being due to differences in precautionary savings behavior.

*abbottbrant@gmail.com or brant.abbott@yale.edu, 778-994-2884 (Canada and USA)

1 Introduction

This paper investigates the effect of being a single-parent on human capital investments in children. There are several potential reasons why a child growing up in a single-parent household might have worse outcomes than in a traditional household, even if total family resources are equal. One possible reason is non-substitutability of inputs into the production of human capital. There is no market for parental time inputs, thus costs will be incurred if neither the time of a single-parent nor any market purchased alternative is a perfect substitute for the absent parent's time. Another possibility is that the imperfections that arise due to incomplete markets for insurance against economic risk might also affect single-parents more than traditional families. A spouse is a very important source of economic insurance (Blundell et al., 2012) that single-parents lack. For this reason single parents are more reliant on self-insurance through precautionary savings, which might crowd out investments in children. The presence of a spouse also implies an advantage for traditional families in handling shocks when they arrive because of intra-household labor supply and child care adjustments. The presence of any of these disadvantages would have important implications for the design of public policy because most progressive programs, for example targeted public preschool, implicitly target single-parent families who are poorer on average.

I develop a heterogeneous agents incomplete markets model that captures the multiple channels through which market incompleteness and missingness affect the outcomes of children. The parameters that determine the disadvantages faced by single-parent households can be identified using within household variation from the PSID Child Development Supplement. The GMM estimates I present indicate imperfect substitutability of parental time and market purchased inputs, as well as imperfect substitutability between parents' time inputs. In the benchmark of the model there is a 37.4% gap between the wages of people who grew up in single-parent versus traditional households. Sixty percent of this is due to differences in parental background and wealth, slightly more than 20% is due to imperfectly substitutable inputs, and slightly less than 20% is due to differences in insurance opportunities.

This paper contributes to the macro-labor literature on childhood skill formation. As outlined in a survey by Heckman and Mosso (2014), most models in this literature have assumed a single-parent single-child structure. Those models that have included multiple parents, e.g. Del Boca et al. (2014), have made arbitrary assumptions about the substitutability between different parents' time inputs and between parental time and market purchased inputs. Thus the estimates of the substitutability parameters provided in this paper will be useful for the broader literature, not just those seeking to understand the disadvantages faced by single-parent households.

The goals of the macro-labor skill formation literature have been (1) to understand what the important inputs for child development are and what the nature of the technology mapping them to outcomes is, and (2) to understand what mechanisms might distort decisions regarding those inputs, which would lead to a role for public policy. The largest contributions to the first goal have been provided by Cunha et al. (2010), Del Boca et al. (2014), and Todd and Wolpin (2007). Papers whose main contributions are more related to public policy include Cunha (2013), Caucutt and Lochner (2012), and Del Boca et al. (2012). Related work has also explored the important of parental beliefs (Carneiro et al., 2013) and the timing of parental income (Carneiro et al., 2014). Finally, Gayle et al. (2013) find that single-parenting rates are the most important factor in determining the parental time input gap between white and black households.

This paper also relates to the small empirical economics literature that explores the effect of growing up with a single-parent. The findings of this literature have been mixed with some papers finding significant detrimental effects (Gruber, 2004; Corak, 2001), and others unable to reject null effects (Lang and Zagorsky, 2001). One of the difficulties has been finding exogenous variation in single-parenting in observational data, and often parental death has been taken as at least part of the exogenous variation. Some of the value of the structural approach I adopt is that it allows me to derive implications for observational data from which the parameters can be identified.

Child development and single-parenting have, of course, been studied in detail by researchers from many disciplines other than economics (see Lamb (2004) for a survey). Economic channels, namely the reduced financial and time resources available in single-parent households, are emphasized in this literature. However, a great deal of importance is also placed on the tendency for single-parents to lack emotional support and develop feelings of isolation that usually would not occur with a spouse.[1] This creates higher stress levels in single-parents, which reduces the effectiveness of parenting, or even causes neglect. Some researchers conclude that the stress experienced by low-socioeconomic status households is the most important determinant of the differences in the outcomes of children (Bradley and Corwyn, 2002). One way to think about this channel within an economic model is that there is complementarity between the time inputs of spouses: each hour of time is less effective when a parent provides care on their own because of the higher stress levels.

The presence of uninsurable economic risk has been one of the suggested motivations for policy interventions in early childhood skill formation. Households wish to self-insure and, because human capital investments are risky, this shifts investments towards risk free assets leading to under investment in children's human capital (See Cunha, 2013).[2] Because single-parents lack the important source of insurance that a spouse represents they will be more reliant on self-insurance, which causes a downward distortion of their child investment decisions. The absence of a spouse also affects the ability of single-parent households to handle shocks when they arrive, which is best illustrated by an example. Suppose the head of household experiences a positive wage shock. In a traditional household the head can increase their labor supply to take advantage, and any reduction in their parenting

[1]This effect is compounded by the fact that the economic disadvantage experienced by single parents creates stress and a greater need for emotional support.

[2]Of course, if human capital investments were riskless then parents would over invest as in the model of Aiyagari et al. (2002).

time can be made up by their spouse. In a single-parent household increasing labor supply to take advantage will mean reduced parenting time, and the only way to compensate is with market purchased inputs, such as daycare. As Appendix A shows, single-parent and traditional households do indeed respond differently to changes in economic circumstances. If market inputs are imperfect substitutes for parenting time then costs will be incurred in order for the parent to take advantage of the positive wage shock. Thus the total effect of economic risk, and the extent to which it differentially affects single-parents, is also related to the substitutability of child development inputs.

Although market purchased inputs can be used as a substitute when the shadow price of parental time increases, they should not simply be viewed as substitutes. Market purchased inputs, such as daycare, music lessons and sports enrollment, are themselves productive inputs that may be complementary to parental time. This is captured through the nested CES human capital investment function I adopt. The outer CES is between market purchased inputs and total effective parental time, where the weight on market inputs is allowed to vary with the age of the child. Total effective parental time depends on an inner CES aggregation of the parents' time inputs, which allows me to test for the importance of participation by both spouses. All of the parameters of this composite investment function are estimated before it is nested within the Cunha et al. (2010) skill formation technology. This technology captures the multi-dimensional self-productive nature of skills, as well as the sensitivity of skills to investment timing. By combining this technology with the composite investment function I am also able to capture the distinction between parental time investments and investments of purchased goods and services, the importance of which was illustrated by Del Boca et al. (2014).

Individual's entire life cycles are modeled. Parents make human capital investment decisions when their children are young. When a child has grown up they are endowed with assets as well as fertility timing and marital preferences. To be consistent with the data, the distributions of these preferences depend on abilities. Next, households are then formed, either consisting of a couple or a single-parent. After a period of time, which depends on fertility timing preferences, children are born. At this point households are motivated by altruism to begin investing in their children. After their children have grown up households live out the rest of their working careers and retirement.

The distributions of cognitive and non-cognitive skills, the productivity of labor, output and inequality are all endogenously determined in the model. This allows me to study the importance of each of the channels that determines the outcome gaps between children from single-parent and traditional households though counterfactual simulations. I am also able to simulate partial equilibrium policy changes, which indicate the steepness of the gradients of economic efficiency and social welfare with respect to policies targeted in various directions. I find that these outcome measures improve as subsidies for investments in children are introduced. I also find that the rate at which these outcome measures increase is greater if the subsidies are targeted to investments by single-parents.

2 A Model of Child Development and Parenting Decisions

2.1 Households

A life is modeled as lasting $J + 17$ years. However, the household makes autonomous decisions only during the J adult years. A model period represents two years of life, thus there are nine child development periods $0 - 1$, $2 - 3$, ..., $14 - 15$, $16 - 17$. Upon reaching 18 years of age people begin making their own autonomous consumption, savings and labor supply decisions. At age j^* children are born into the household, and the scope of decisions the new parent must make expands to include time and expenditure inputs for their children's development. When the children become adults their parents, who will be j^{**} years old, earn an altruistic payoff based on the children's expected lifetime utility. This altruistic payoff is what motivates parents to invest in their child. After the children have become autonomous their parents carry on working until retirement at age, j^{RET}, when they begin to consume out of savings and social security. During retirement there is mortality risk, and the probability of surviving to age $j + 1$ conditional on living to age j is Ψ_{j+1}. Survival probabilities are set using the US Life Tables, 2007.

The model distinguishes cognitive and non-cognitive skills, θ_C and θ_N, respectively. For notational ease these are grouped into a skill vector θ. A child's skill vector is denoted $\hat{\theta}$ to distinguish it from their parents. The complexity of the model does not permit college education decisions to be considered formally; however, the cost of college is an important consideration for families. Thus I include college costs for high achieving children by assuming that any child for whom $\hat{\theta} >> \underline{\theta}$ at age 18 attends college. I calibrate the $\underline{\theta}$ threshold so that the top 35% of the ability distribution graduates from college.

Marital preferences are determined exogenously. Upon reaching adulthood a person draws χ from a binomial distribution. If $\chi = 2$ then the person will be married, and if $\chi = 1$ they will be single. The probability of being married is a function of abilities $\pi(\theta)$ in order to capture the positive correlation between ability and marital status observed in the data.[3] I assume perfect assortative mating on skills. Once formed, families draw fertility timing preferences μ. The distribution of timing preferences $f(\mu|\theta)$ is also conditional on skills. I assume that numbers of children per household are the median values observed in Census data, which are one child for single-parents and two children for married couples. In the quantitative exercise I assume that married households are collective. All of the strong assumptions that have been discussed in this paragraph apply only to the quantitative model and are needed for tractability. The estimation is much more flexible and does not require these assumptions, nor does it require the specific utility functions specified next.

[3] Below I document the very strong positive correlation between AFQT scores and marital status observed in NLSY79 data.

2.1.1 Preferences and Altruism

Periodic preferences of an individual depend on leisure $\ell_{i,j}$ and consumption $c_{i,j}$. Utility is also time separable and future payoffs are discounted by a factor β per period. A weight ω is applied to the altruistic payoff one attains through each of their children, which is the child's expected value function $\Omega(\hat{a}, \hat{\theta})$, evaluated at the child's final human capital attainment, $\hat{\theta}$, and wealth endowment, $\hat{a}$.

For a two-parent household lifetime utility is

$$\sum_{j=1}^{J} \beta^{j-1}[U(c_{1,j}, \ell_{1,j}) + U(c_{2,j}, \ell_{2,j})] + 2\beta^{j^{**}-1}\omega\Omega(\hat{\theta}, \hat{a}). \tag{1}$$

The altruistic part of preferences is multiplied by two because two-parent households have two children. For a single-parent household lifetime utility is similar except there is only one child and one adult.

$$\sum_{j=1}^{J} \beta^{j-1}U(c_{1,j}, \ell_{1,j}) + \beta^{j^{**}-1}\omega\Omega(\hat{\theta}, \hat{a}). \tag{2}$$

Note that the heterogeneity in fertility timing preferences means that j^{**} is implicitly a function of μ.

In the quantitative model utility functions are assumed to be separable:

$$U(c_i, \ell_i) = \frac{c_i^{1-\sigma}}{1-\sigma} + \nu\frac{\ell_i^{1-\frac{1}{\zeta}}}{1-\frac{1}{\zeta}}. \tag{3}$$

The (two-year) discount factor is set to $\beta = 0.95$, the intertemporal risk aversion parameter is set to $\sigma = 2$ and the Frisch elasticity of labor supply is set to $\eta = 1/2$.

2.1.2 Child Development Technology

A child's age is indexed by k. Throughout the periods in which a parent cares for their child they can improve the child's human capital through a technology that maps the child's existing skills, $\hat{\theta}_k$, time investments, $t_{1,k}$ and $t_{2,k}$, expenditure investments, x_k, parental abilities, θ, and development shocks, η_k, into the child's future skills, $\hat{\theta}_{k+1}$. In general the skill formation technology is a function g:

$$\hat{\theta}_{k+1} = g(k, \hat{\theta}_k, \theta, x_k, t_{1,k}, t_{2,k}, \eta_k). \tag{4}$$

The implementation of g follows Cunha et al. (2010), where the aggregation of time and goods investments into a single index is represented by $\mathcal{I}$:

$$\hat{\theta}_{s,k} = \exp(\eta_{s,k}) \left[\gamma_{1,s,k}\hat{\theta}_{C,k}^{\psi_{s,k}} + \gamma_{2,s,k}\hat{\theta}_{N,k}^{\psi_{s,k}} + \gamma_{3,s,k}\mathcal{I}_k^{\psi_{s,k}} + \gamma_{4,s,k}\theta_C^{\psi_{s,k}} + \gamma_{5,s,k}\theta_N^{\psi_{s,k}}\right]^{\frac{1}{\psi_{s,k}}} \tag{5}$$

Cognitive and non-cognitive skills are indexed by $s \in \{C, N\}$.

Estimates of the sets of γ and ψ parameters, as well as the variances of the set of η shocks, are available for two age groups in Cunha et al. (2010). I specifically employ estimates in Table 4 of that work, which reports estimates when unobserved heterogeneity has been accounted for (although their results are all quite robust). Estimation in that paper is performed for two age groups of children, but rather than have the parameters jump I assume a smooth linear transition across the development years. At child ages 0-1 each parameter takes the first-stage estimate from Cunha et al. (2010), and at ages 16-17 each parameter takes the second-stage estimate. The parameters evolve linearly during the intermediate years.

To accommodate investments of goods and each parents time the investment index in equation (5) is modeled as a nested CES function of expenditure and time inputs:

$$\mathcal{I}_k = \delta \left(\alpha_k x_k^{\lambda_0} + (1-\alpha_k)[t_{1,k}^{\lambda_1} + t_{2,k}^{\lambda_1}]^{\frac{\lambda_0}{\lambda_1}} \right)^{\frac{1}{\lambda_0}} . \tag{6}$$

The outer CES function captures the potential non-substitutability between time and goods investments through the parameter λ_0, and also captures the possibility of an increasing importance of expenditure investments as children get older through age varying α_k. The inner CES captures the possibility of complementarity between the time inputs of parents. If $\lambda_1 = 1$ then there is no cost to a single-parent providing all of the time inputs. However, if $\lambda_1 < 1$ there is a gain in total effective time when the inputs are spread between two parents, one reason for which might be the stress associated with single-parenting.

2.1.3 Wages, Time Endowments and Human Capital

Each adult is endowed with a unit of time, which they allocate between labor supply, leisure and, when applicable, investing in their child. Wages paid per unit of labor supply depend on cognitive and non-cognitive skills, an age-experience profile, ξ_j, and the idiosyncratic productivity shock, z_i:

$$\ln W(\theta, z_i, j) = \kappa_C \left(\frac{\ln \theta_C - \overline{\theta}_C}{\sigma_C} \right) + \kappa_N \left(\frac{\ln \theta_N - \overline{\theta}_N}{\sigma_N} \right) + \xi_j + z_i, \tag{7}$$

where i can be either the head or spouse. Notice that the factor loadings on abilities, κ_s, differ for cognitive and non-cognitive skills. The idiosyncratic productivity component is modeled as a persistent mean zero process

$$z_i' = \rho z_i + v_i', \tag{8}$$

where v_i' is an iid normal process.

The wage process is parameterized based on a combination of pooled estimates of the process estimated in Abbott et al. (2013), and the ability factor loadings estimated by Heckman et al. (2006). The persistence and volatility of estimated wage shocks is consistent with other literature, e.g. Storesletten et al. (2004) or Heathcote et al. (2010). Specifically, the parameters the z_i process in equation (8) - adjusted for the two-year length of the model periods - are $\rho = 0.92$, $\sigma^2_{z0} = 0.12$ and $\sigma^2_v = 0.04$. The age process is a quartic fit to PSID data (see Abbott et al. (2013) for parameter values), and the loading factors on cognitive and non-cognitive abilities are $\kappa_C = 0.22$ and $\kappa_N = 0.046$ (see Heckman et al. (2006)).[4]

2.1.4 Assets and Annuities

There is a single risk free asset in the economy. The quantity owned by a household at the beginning of age j is a_j. This asset pays a return r if positive quantities are owned by the household, and households who hold negative quantities of the asset (debt) must pay interest, also at the rate r. A natural borrowing limit is enforced, which ensures that all debts are repayable. The positive aggregate holdings of this asset are invested in physical capital, which also earns a net return r. The exogenous (biennial) interest rate is set to $r = 0.05$.

An income based credit constraint is imposed on working age households. They may borrow up to the equivalent of 74% of quarterly household income, which Kaplan and Violante (2014) find is the median value in the Survey of Consumer Finances data. Perfect annuity markets are assumed during retirement. The wealth of the $1 - \Psi_j$ mass of retired households who die between ages $j-1$ and j is redistributed to the surviving members of that generation.

2.1.5 Taxes and Transfers

I implement a nonlinear tax and transfer function, which depends on marital status and children. The functions $T_\chi(y)$ apply to households without children, and the functions $\widetilde{T}_\chi(y)$ apply to households with children (recall that $\chi \in \{1, 2\}$ denotes single or married). Total taxable income y includes labor and capital income. These functions are parametric representations of the progressive income tax system, the Earned Income Tax Credit, the Child Tax Credit, food stamps (SNAP), and other progressive transfers. The foundation for these functions is the work by Guner et al. (2014), which considers income taxes and EITC. To this I add food stamps for the poorest households according to the SNAP guidelines, which include and implicit 30% income tax, and the Child Tax Credit, which is a refundable $1000 per child credit. I add an additional basic transfer received by all households in order to replicate the overall progressivity of the tax system, and the fact that even very high income households often get transfers.[5]

[4] Heckman et al. (2006) differentiate gender in their estimates, although the point estimates are similar. The parameters adopted here are midpoints between male and female estimates.

[5] Even households earning a quarter of a million dollars receive $6000 in transfers annually on average: http://www.cbo.gov/sites/default/files/cbofiles/attachments/43373-06-11-HouseholdIncomeandFedTaxes.pdf.

2.1.6 Equivalence Scales

One of the commonly cited sources of marital surplus is economies of scale where, for example, family members share housing resources (e.g. Browning et al., 2013).[6] To capture this I include household equivalence scales $o(j, \mu, \chi)$, which are based on the mean benchmark scale developed by Fernandez-Villaverde and Krueger (2007). For a single-parent family during periods that children are not in the household $o(j, \mu, \chi) = 1$. For the remaining periods $o(j, \mu, \chi) < 1$ denotes the inverse of the Fernandez-Villaverde and Krueger (2007) mean scale.

2.2 Decision Problems

The decision problems are presented in recursively. Stages that have unique decision problems are work before children are born, the child development periods, the period when the child becomes independent, work after the child moves out and retirement. In periods when children are absent from the households the value function for household type χ is denoted V_χ. When the household has children the value function is $\widetilde{V}_\chi$, and when the household is retire the value function is V_χ^R.

2.2.1 Retirement

During retirement each individual receives pension income that depends on their skills $p(\theta)$. Skills are the largest determinant of average lifetime income thus this approach closely approximates the true social security system. The optimization problem of a single retired household is given in equation (9).

$$\begin{aligned} V_1^R(j, a, \theta) &= \max_{\{c_1, a'\}} \left\{ U(c_1, 1) + \Psi_{j+1} \beta \mathbb{E} \left[V_1^R(j+1, a', \theta) \right] \right\} \qquad (9) \\ s.t. & \\ c &= \frac{1}{\Psi_j}(1+r)a - a' + p(\theta) - T_1(y) \\ a' &\geq 0 \end{aligned}$$

Note that the inflation of assets by $1/\Psi_j$ reflects the perfect annuity markets assumption.

For a married retired household the value function V_2^R would differ from equation (9) in the following ways. Flow utility would be the sum across both members payoffs, there would bet two pension incomes, the tax function would be that for a two member household, and consumption would be multiplied by the inverse equivalence scale $o(j, \mu, \chi)$, which equals unity for single-person retired households.

[6] Another issue is how resources are divided up within the household, which can be identified under some assumptions (Dunbar et al., 2013). Here I ignore this issue and focus on the total utility that consumption brings to the family.

The function $p(\theta)$ is set as $p \cdot \theta_c$, which is linear in cognitive skills. I calibrate p so that the average replacement rate in the economy is 45%.

2.2.2 Work After the Child's Independence

From age $j^{**}+1$ until retirement households solve a stochastic consumption savings problem. For a traditional household this problem is:

$$\begin{aligned}
V_2(j,a,z_1,z_2,\theta,\mu) &= \max_{\{\ell_1,\ell_2,c_1,c_2,a'\}} \left\{ \sum_{i=1}^{2} U(c_i,\ell_i) + \beta \mathbb{E}\left[V_2(j+1,a',z_1',z_2',\theta,\mu)\right] \right\} \\
s.t. & \\
o(j,\mu,\chi)\,(c_1+c_2) &= (1+r)a - a' + \\
& \quad + W(\theta,z_1,j)(1-\ell_1) + W(\theta,z_2,j)(1-\ell_2) - T_2(y) \\
a' &\geq \underline{a} \times \left[W(\theta,z_1,j)(1-\ell_1) + W(\theta,z_2,j)(1-\ell_2)\right] \\
z_i' &= \rho z_i + v_i'
\end{aligned} \tag{10}$$

The variable y represents the total taxable income of the family.[7] The amount they can borrow depends on the combined incomes of both the head and spouse. For a single-person household this problem would be altered to include only one person's flow utility, one person's decision variables and one productivity shock. The tax function would also change to that for a single person household, and the equivalence scale would equal unity.

2.2.3 Children Move Out

At age j^{**} a family's children turn 18 and move out to form their own households. When they do their parents earn an altruistic payoff based on their expected lifetime utility. At this point parents can make a financial transfer to their children. State variables include the skills, $\hat{\theta}$, and age, $k=18$, of the child at this stage. For a single-parent the problem solved at this stage is

$$\begin{aligned}
\widetilde{V}_1(j,k,\hat{\theta},a,z_1,\theta,\mu) &= \max_{\{\hat{a},\ell_1,c_1,a'\}} \left\{ U(c_1,\ell_1) + \beta \mathbb{E}\left[V_1(j+1,a',z_1',\theta,\mu)\right] + \omega\Omega(\hat{a},\hat{\theta}) \right\} \\
s.t. & \\
c_1 + \hat{a} &= (1+r)a - a' + W(\theta,z_1,j)(1-\ell_1) - T_1(y) \\
a' &\geq \underline{a} \times \left[W(\theta,z_1,j)(1-\ell_1)\right] \\
\hat{a} &\geq 0 \\
z_1' &= \rho z_1 + v_1'.
\end{aligned} \tag{11}$$

[7] Defined as labor plus capital income.

The timing is such that the child leaves the household at the beginning of the period with the transfer in hand. One implication is that the parent's taxes no longer include benefits from having children in the household.

In the traditional household version of this problem the altruistic payoff is $2\omega\Omega(\hat{a},\hat{\theta})$, reflecting the fact that traditional households have two children each. Furthermore, the cost of transferring $\hat{a}$ to each child would be reflected in the budget constraint by $2\hat{a}$ entering on the left-hand-side. Indeed, traditional households optimally treat their children equally in this model, and hence the children are identical at the time they move out. The tax function for two-adult households and the consumption equivalence scale also enter the traditional household version of the problem.

A child's expected lifetime utility is

$$\Omega(\hat{a},\hat{\theta}) = \int \left(\pi(\hat{\theta})\mathbb{E}[V_1(1,\hat{a},z_1,\hat{\theta},\mu)] + (1-\pi(\hat{\theta}))\frac{1}{2}\mathbb{E}[V_2(1,2\hat{a},z_1,z_2,\hat{\theta},\mu)]\right) f(\mu|\hat{\theta})d\mu. \quad (12)$$

Expectations are taken over the distributions of fertility timing preferences $f(\mu|\hat{\theta})$ and marital status $\pi(\hat{\theta})$. I assume that parents value half of the expected utility generated by the marriage of their child. Perfect assortative mating means that in the event of marriage each child will bring $\hat{a}$ assets to the new household.

An important aspect of this problem is the constraint $\hat{a} \geq 0$, which enforces the assumption that parents cannot borrow against their child's future income (Laitner, 1992). An important implication of this is the inability of parents to optimally invest in a child's human capital using the child's future earnings, even if the child would want them to.

2.2.4 Child Development Periods

From age j^* to $j^{**}-1$ a parent invests in their child, in addition to consuming and working themselves. Children's current skills and age are state variables at this stage because they are arguments of the skill formation technology g, which was described in equations (5) and (6) above. A single-parent family solves the following problem:

$$\begin{aligned}
\widetilde{V}_1(j,k,\hat{\theta},a,z_1,\theta,\mu) &= \max_{\{x,t_1,\ell_1,c_1,a'\}} \left\{U(c_1,\ell_1) + \beta\mathbb{E}\left[\widetilde{V}_1(j+1,k+1,\hat{\theta}',a',z_1',\theta,\mu)\right]\right\} \\
s.t. & \\
o(j,\mu,\chi)c_1 + x &= (1+r)a - a' + W(\theta,z_1,j)(1-\ell_1-t_1) - \widetilde{T}_1(y) \\
\hat{\theta}' &= g(k,\hat{\theta},\theta,x,t_1,\bar{t}_2,\eta) \\
a' &\geq \underline{a} \times W(\theta,z_1,j)(1-\ell_1-t_1) \\
\eta &\sim N(0,\Sigma_{\eta,k}) \\
z_1' &\sim \rho z_1 + v_1'
\end{aligned} \quad (13)$$

Like assets, children's skill vectors are endogenous state variables resulting from current investment decisions. By improving skills now the parents anticipate a greater altruistic payoff in the future. One important point is that t_2, the time input of the absent parent, is not a choice variable. In the data absent spouses spend a small amount of time with their children, and I capture this in the model by exogenously setting $t_2 = \bar{t}_2$, which is the sample average of absent parent time supply reported in Table 2.

For traditional households the decisions problem during the child development periods is

$$
\begin{aligned}
\widetilde{V}_2(j,k,\hat{\theta},a,z_1,z_2,\theta,\mu) &= \max_{\{x,t_1,t_2,\ell_1,\ell_2,c_1,c_2,a'\}} \left\{ \sum_{i=1}^{2} U(c_i,\ell_i) + \right. \\
&\quad \left. +\beta\mathbb{E}\left[\widetilde{V}_2(j+1,k+1,\hat{\theta}',a',z_1',z_2',\theta,\mu)\right] \right\} \\
s.t. & \\
o(j,\mu,\chi)(c_1+c_2)+\phi_x x &= (1+r)a - a' + W(\theta,z_1,j)(1-\ell_1-\phi_t t_1) \\
&\quad + W(\theta,z_2,j)(1-\ell_2-\phi_t t_2) - \widetilde{T}_2(y) \\
\hat{\theta}' &= g(k,\hat{\theta},\theta,x,t_1,t_2,\eta) \\
a' &\geq \underline{a} \times [W(\theta,z_1,j)(1-\ell_1-\phi_t t_1) + W(\theta,z_2,j)(1-\ell_2-\phi_t t_2)] \\
\eta &\sim N(0,\sigma^2_{\eta,k}) \\
z_i' &\sim \rho z_i + v_i'.
\end{aligned} \tag{14}
$$

One of the important differences to note is the possibility of economies of scale in investments, which are captured by ϕ_x and ϕ_t. If, for example, there are economies of scale in time investments then providing one unit of parental time to each child costs $\phi_t < 2$ units of the respective parent's time endowment. If there are no economies of scale then $\phi_t = 2$.

The periodic skill development shocks, η, are assumed to occur at the end of the period, after investments have already been decided. This assumption ensures that the model is consistent with the empirical formulation in Cunha et al. (2010), where the shocks are assumed to be orthogonal to inputs.[8] To keep the analysis tractable I require that the skill vectors of siblings are identical, thus I assume that siblings experience equivalent development shocks.

2.2.5 Work Before Children

Prior to the birth of their children households usually solve the same problem as after their children move out, the traditional household version of which is given in equation (10). There are two exceptions to this: college years for high achieving individuals, and the period just before children are

[8]Unobserved heterogeneity in the productiveness of inputs, as well as measurement error, are separately accounted for in their formulation.

born.

In the year prior to the child being born a single-parent household's continuation value is

$$\mathbb{E}_{z_1'}[\widetilde{V}_1(j^*, 0, \hat{\theta}_0, a', z_1', \theta, \mu)]. \tag{15}$$

I assume that all children begin with the same stock of skills, $\hat{\theta}_0$, which is the lowest point on the ability grid.[9,10] The zero in the second argument simply indicates that newborns are zero periods old.

For those individuals who go to college the associated costs are incurred during the first two periods (four years) of adulthood. These costs include time to study, which is assumed to be 1/3 of the time endowment, and tuition, which is assumed to be the model equivalent of $6000 per year. Other than the additional debits from the time and budget constraints college students solve the same problem as non-college workers who have no children currently.

3 PSID Child Development Supplement Data

The Child Development Supplement of the Panel Study of Income Dynamics (PSID-CDS) began collecting data in 1997 with the goal "to provide researchers with a comprehensive, nationally representative, and longitudinal database of children and their families with which to study the dynamic process of early human capital formation." The initial sample included 3563 children under 13 years old whose parents were included in the main PSID sample. Up to two children per household were included. Second and third waves of data collection were carried out in 2002-03 and 2007. These additional waves re-interviewed 2907 and 1506 eligible children under the age of 19.

The gem of the PSID-CDS is the time-use survey, which provides comprehensive information on the activities of each child, and the inclusion of others in those activities. The time diaries include 24 hours of information for two representative days, one a weekday and the other a weekend day. For each activity a child does during survey day information is collected on the duration of the activity, who was participating with the child, who was around but not participating, and where the activity took place. I define the time inputs of parents as the total duration of activities in which they are actively participating in the activity the child is engaged in. An estimate of weekly totals is based on five times the weekday time allocation plus two times the weekend time allocation. Note that even if a parent is absent from the household their time inputs are still observed if the child spends time with them.

Observed child-specific expenditures were limited in 1997 and appear unreasonably low; however,

[9] I make this assumption because any parameterization of the link between initial skills and *in utero* conditions would be tenuous.

[10] Recent research has demonstrated the importance of pre-birth factors, e.g. Hoynes et al. (2012). However, these factors are part of parental background, and the emphasis here is on differences across households other than those driven by background.

in the 2002 and 2007 waves many additional expenditure items were added resulting in a more reasonable total expenditure variable. Items include tuition costs, tutoring expenses, lessons (i.e. music lessons), sports, community/religious groups, toys/books, vacations, school supplies, clothing, transportation and daycare expenses. The 2002 wave also includes spending on food specifically for the child, but this item was not included in 2007 and I exclude it in 2002 as well for consistency.

Several sub-tests of the Woodcock-Johnson Achievement Test and Wechsler Intelligence Scale for Children were administered to children of various ages. Like in Del Boca et al. (2014), I use raw vocabulary Letter-Word (LW) scores as a measure of cognitive ability, and additionally use the math Applied-Problems (AP) raw scores. All three of these tests have good breadth of coverage in the 2002 and 2007 samples. The LW test involves oral identification of a sequence of 57 words that become progressively more difficult, although the first few items are simple 'letter recognition' tests. Success on an item will require the child to either know the word, and hence state it, or to have strong enough reading skills to pronounce an unknown word. The AP test score is based on the number of math reasoning problems correctly answered. The AP problems are intended to measure reasoning ability, thus calculations are relatively simple if the correct procedure is identified.

The subset of the PSID-CDS employed in the current study are those children for whom complete ability, time-use and child-expenditure data are observed, and furthermore parental responses must be observed in the 2003 and 2007 PSID main data. The 2003 main PSID data were chosen as the counterpart to the 2002 CDS because the 2002 CDS was actually collected over the 2002-03 period. When estimating the model the parental wages need to be observed in both 2003 and 2007 (depending on the equation). The observations that are missing wage information in either time period are still included as part of the sample because they are used to estimate the selection equations. Observations are dropped if household composition, taxable income or state are missing because these are needed in order to use the NBER TAXSIM software to compute net wages.

It is useful to provide several snapshots of the data in which they are partitioned in different ways. The first is provided in Table 1, which present summary statistics for the pooled 2002/2007 waves by the age of the child. These data include all observations even those that will eventually be excluded due to unobserved wages. The most notable patterns are the decreasing time inputs and increasing expenditure inputs as the children get older. Naturally, test scores increase as the children get older.

Next, Table 2 focuses on the topic of this paper, which is differences across single-parent and traditional households. This table uses data only from the 2002 wave of the CDS and excludes (for clarity) the 5% of single parents who are fathers. The general disadvantage of children in single-parent households in terms of total resources invested is evident in these statistics. The time input of a mother is nearly identical whether she is single or married, but the time input of the father is less and one-quarter as large if he is absent. Investments of goods and services are also scarcer for children of single parents, being just over 20% smaller on average. Single mothers work 500 hours per year more than married mothers even though they spend an equal amount of time with their children.

Lastly, Table 3 illustrates the variation across waves (time periods) in the data that are used to estimate the model. Only about 2/3 of observations (543 children) have non-missing wage data for the head across both time periods. Note that changes in marital status usually result in a change of the head of households. Any observations with a change in who the head or wife is are dropped because the observed wages in the second period are not for the same person as the original observation.

4 Estimation of Child Investment Parameters

4.1 Theory Based Moment Conditions

The parameters to be estimated are those of the composite investment function

$$\mathcal{I}_k = \delta \left(\alpha_k x_k^{\lambda_0} + (1-\alpha_k)[t_{1,k}^{\lambda_1} + t_{2,k}^{\lambda_1}]^{\frac{\lambda_0}{\lambda_1}} \right)^{\frac{1}{\lambda_0}} . \tag{16}$$

Specifically, I estimate the age-varying expenditure weights α_k, the substitutability of time and goods λ_0, and the substitutability between the time inputs of the head and the spouse λ_1. The scale parameter δ will be determined in the model calibration as it will not be identified in estimation.

The outcome variable $\mathcal{I}_k$ measures a composite investment that is not observed. Thus, I will adopt a marginal values approach that relies on implications of household optimization given prices (a seminal example of this approach is Heckman, 1974). This approach is feasible here because the inputs are observed, and in most cases the relative prices (wages) are observed (selection issues are addressed below).

The theory provides three first-order conditions for the optimal input decisions that will be used to identify the parameters of interest. These reflect the optimal inputs of expenditure x_k, the head's time $t_{1,k}$, and the spouses time $t_{2,k}$:

$$\begin{aligned}
\beta V_{\mathcal{I}} \frac{\partial \mathcal{I}_k}{\partial x_k} &= \beta V_{\mathcal{I}} \delta \mathcal{I}_k^{1-\lambda_0} \alpha_k x_k^{\lambda_0 - 1} = \phi_x U_c & (17)\\
\beta V_{\mathcal{I}} \frac{\partial \mathcal{I}_k}{\partial t_{1,k}} &= \beta V_{\mathcal{I}} \delta \mathcal{I}_k^{1-\lambda_0} (1-\alpha_k)[t_{1,k}^{\lambda_1} + t_{2,k}^{\lambda_1}]^{\frac{\lambda_0}{\lambda_1}-1} t_{1,k}^{\lambda_1 - 1} = \phi_t U_{\ell_1} & (18)\\
\beta V_{\mathcal{I}} \frac{\partial \mathcal{I}_k}{\partial t_{2,k}} &= \beta V_{\mathcal{I}} \delta \mathcal{I}_k^{1-\lambda_0} (1-\alpha_k)[t_{1,k}^{\lambda_1} + t_{2,k}^{\lambda_1}]^{\frac{\lambda_0}{\lambda_1}-1} t_{2,k}^{\lambda_1 - 1} = \phi_t U_{\ell_2}, & (19)
\end{aligned}$$

where $V_{\mathcal{I}}$ is the inner product between the gradient of the value function on skills and the gradient of skills on investment:

$$V_{\mathcal{I}} = \mathbb{E}\left[\frac{\partial \tilde{V}_2(j+1, k+1, \hat{\theta}', a', z_1', z_2', \theta, \mu)}{\partial \hat{\theta}'} \cdot \frac{\partial \hat{\theta}'}{\partial \mathcal{I}_k} \right] . \tag{20}$$

Note that for single-parent families equation (19) does not apply because $t_{2,k}$ is not a choice when no spouse exists, and the scale parameters ϕ_x and ϕ_t are unity. These differences are taken into account below. The value function in equation (20) would also be replaced by that of a single-parent family.

Dividing equations (18) and (19) each by equation (17), and then further dividing (19) by (18) yields more useful forms for these conditions:

$$\frac{1-\alpha_k}{\alpha_k} x_k^{1-\lambda_0}[t_{1,k}^{\lambda_1}+t_{2,k}^{\lambda_1}]^{\frac{\lambda_0}{\lambda_1}-1} t_{1,k}^{\lambda_1-1} = \frac{\phi_t U_{\ell_1}}{\phi_x U_c} \tag{21}$$

$$\left(\frac{t_{2,k}}{t_{1,k}}\right)^{\lambda_1-1} = \frac{U_{\ell_2}}{U_{\ell_1}}. \tag{22}$$

The marginal utilities are unobservable, therefore the right-hand sides of these equations will need to be replaced. To do this the relationship between wages and the marginal utilities from consumption and leisure will be used. Whenever an adult member of the household works the usual optimal labor supply condition holds:

$$\frac{U_{\ell_i}}{U_c} = (1-T'(y_i))W_i = \tilde{w}_i. \tag{23}$$

This restriction is sufficient to replace the right-hand-side terms in equations (21) and (22) with observable wages.

It is both noteworthy and important that far less needs to be assumed about the utility functions in order to derive these equations than was assumed for tractability of the main quantitative framework. Utility functions with complementarity across spouse's leisures or between consumption and leisure are allowed for. Furthermore, it is not necessary to assume collective households to identify the parameters. With collective households spousal marginal utilities from consumption cancel out of the second equation, but here I leave them in to allow for limited commitment or other non-cooperative bargaining arrangements.

After substituting wages into equations (21) and (22), taking logs and introducing error terms, I the arrive at the following data counterpart for household i to the theoretical moment equations:

$$\begin{aligned} \ln(\frac{1-\alpha_k}{\alpha_k}) + (1-\lambda_0)\ln(x_{k,i}) + (\frac{\lambda_0}{\lambda_1}-1)\ln(t_{1,k,i}^{\lambda_1}+t_{2,k,i}^{\lambda_1}) + \\ +(\lambda_1-1)\ln(t_{1,k,i}) - \ln(\phi_t/\phi_x) - \ln(\tilde{w}_{1,i}) &= m_{1,i} \qquad (24) \\ (\lambda_1-1)\ln(t_{2,k,i}/t_{1,k,i}) - \ln(U_{c_2,i}/U_{c_1,i}) - \ln(\tilde{w}_{2,i}/\tilde{w}_{1,i}) &= m_{2,i}, \qquad (25) \end{aligned}$$

where $m_{1,i}$ and $m_{2,i}$ are the error terms.

4.2 Unobserved Heterogeneity and Selection

There are two possible sources of selection bias: the first is the usual problem of wage censoring, and the second is the problem that one of the moment conditions can only be observed for intact households. As the following paragraphs describe, both of these issues will be dealt with by assuming that the troublesome heterogeneity is permanent.[11]

Selection into marriage is an issue because equation (25) can only be observed if the parents are either married or cohabitate, otherwise the wage of the absent parent is unobserved. We often observe that absent parents still provide positive time inputs, and if those absent parents also work then the second equation would still apply. What is troubling is that there might be unobserved heterogeneity that affects both the likelihood of marriage and parenting inputs. A simple example would be heterogeneity in preferences for children. Those who like children might be more likely to be married and to also spend extra time with their children. Thus estimating based on married observations only would bias the results relative to estimating based on all observations for whom the first-order condition holds.

To address this I will assume that selection into marriage is based on permanent heterogeneity that drops out in first-differences. That is, while I allow that $m_{2,i}$ is correlated with the probability of marriage/cohabitation, I assume that $\Delta m_{2,i}$ and the probability of marriage/cohabitation are uncorrelated. Changes in marital status can still occur, but the shocks that cause these changes must be uncorrelated with unobserved heterogeneity that enters $m_{2,i}$. Similarly, I assume that the correlation between $\Delta m_{1,i}$ and the probability of marriage/cohabitation is zero.

It is also problematic that the equations are only observed if the necessary wage data is observed. For the first equation only the wage of the head needs to be observed, while for the second both the head and wife must be in the labor force. The classic issue of unobserved heterogeneity being correlated with labor force participation decisions would bias estimation of equations (24) and (25) directly. Like in Keane et al. (1988), the standing assumption will be that only the permanent component of the troublesome unobserved heterogeneity is correlated with heterogeneity in labor force participation decisions. I will also test the robustness of the results to this assumption by estimating an additional specification that includes selection correction terms based on inverse Mill's ratios. Appendix B describes the formulation of these correction terms.

4.3 Measurement Error

There is good reason to believe that the inputs are measured with error. For example, the time input variables are based on one weekday and one weekend-day, and it is unlikely that multiplying these observations by the number of week/end days in a year will equal the total yearly input. These measurement errors will potentially bias the estimated parameters, making the inputs appear more substitutable than they really are.

Two address this I project the observed measurements on a number of instruments in a first stage and use the common signal component in the main estimation. Some obvious instruments are avail-

[11] I will also estimate a specification that includes an external selection equation, and the results are entirely robust.

able in the 'passive' time inputs of parents. Recall that the measured time inputs $t_{i,k}$ are 'active' time inputs, the distinction being that the parent is actually participating in an activity with the child. Additionally, lagged test scores are used as instruments. The logic is that test scores will be positively correlated with the inputs, but orthogonal to the error terms $m_{.,i}$. This orthogonality follows from the theory: any variable in the households current information set, which includes lagged test scores, is orthogonal to the first-order conditions.

Specifically, I regress an input, e.g. $t_{1,k}$, on the passive time inputs of both the mother and father and vocabulary (Letter-Word) and math (Applied Problems) scores. The F-statistic for joint significance is always large enough to easily reject the null, and the R^2 are between 0.15 and 0.24. The fitted values, e.g. $\hat{t}_{1,k}$, are used in the main estimation.

4.4 Identifying Assumptions

Given the solutions to selection, unobserved heterogeneity and measurement error described above, the formal estimating equations can now be written:

$$\begin{aligned} \Delta \ln(\frac{1-\alpha_k}{\alpha_k}) + (1-\lambda_0)\Delta \ln(\hat{x}_{k,i}) + (\frac{\lambda_0}{\lambda_1} - 1)\Delta \ln(\hat{t}^{\lambda_1}_{1,k,i} + \hat{t}^{\lambda_1}_{2,k,i}) + & \\ +(\lambda_1 - 1)\Delta \ln(\hat{t}_{1,k,i}) - \Delta \ln(\tilde{w}_{1,i}) &= \epsilon_{1,i} \quad (26) \\ (\lambda_1 - 1)\Delta \ln\left(\frac{\hat{t}_{2,k,i}}{\hat{t}_{1,k,i}}\right) - \Delta \ln\left(\frac{U_{c_2,i}}{U_{c_1,i}}\right) - \Delta \ln\left(\frac{\tilde{w}_{2,i}}{\tilde{w}_{1,i}}\right) &= \epsilon_{2,i}, \quad (27) \end{aligned}$$

where $\epsilon_{.,i} = \Delta m_{.,i}$. To make estimation feasible I assume that $(1-\alpha_k)/\alpha_k$ grows at a constant rate $a = \Delta \ln((1-\alpha_k)/\alpha_k)$, which has also been assumed in previous research, e.g. Del Boca et al. (2013). Next, only the cross-sectional average change in the ratio of marginal utilities from consumption can be identified. Denote this term by $c = E[\Delta \ln(U_{c_2,i}/U_{c_1,i})]$. If households are collective or fully committed then this term will be zero. A non-zero result could indicate that non-collective households have responded in similar ways to an aggregate shock, for example a construction boom that disproportionately increases men's wages. However, this also means that heterogeneity in this term is swept into the error. This will not be a problem because the identifying assumption underlying the estimator will allow second period outcome variables to be correlated with the error term.[12]

To proceed with developing the identifying assumption, rewrite the estimating equations in vector form as

$$F(X_i, X_i', \Lambda) = \epsilon_i, \quad (28)$$

where $\Lambda = (\lambda_0, \lambda_1, a, c)$, X_i is a matrix containing all observations of each variable in the first time

[12]If households are not fully committed then shocks can result in renegotiation of the intra-household sharing rule. However, these shocks must be unexpected and uncorrelated with data from the first observed period. Thus it is safe to identify the parameters using the assumption that first period data is orthogonal to the error term.

period and X_i' is the second period equivalent of X_i.

I have already discussed that heterogeneity of changes in the intra-household bargaining rule will enter the error term. This is one reason why second period data should not be assumed orthogonal to the error term, but there is an additional reason that would apply even if all households are collective. This additional issue arises from the fact that the intra-temporal first-order conditions are part of a broader stochastic dynamic system of equations that describe optimal household behavior. Expectational difference equations are part of this system and forecast heterogeneity, which is unobservable to the econometrician, is generally part of the unobserved heterogeneity. Because second period variables are correlated with forecast heterogeneity, they are also correlated with the error terms and cannot be included in the GMM 'instruments'.

These endogeneity issues preclude the use of (non-linear) least squares, which would have had the benefit of clarity as to the way the parameters are identified. However, much of this clarity can be maintained without assuming future outcomes are orthogonal to the error terms if one makes the similar identifying assumption that

$$E\left[E\left[\left.\frac{\partial F(X_i, X_i', \Lambda)}{\partial \Lambda}\right| X_i\right] \cdot \epsilon\right] = 0. \tag{29}$$

This assumption says that forecasts of the gradient vector are orthogonal to the error term, whereas least-squares would say that the gradient itself is orthogonal. This type of assumption is similar to the Arellano and Bond (1991) approach to estimating dynamic panel-data models with lagged dependent variables. Framed as a GMM estimator, this assumption results in an exactly identified system of moment conditions, and is valid because the instruments are functions of data that are included in the first-period information set. Note that $E[F_\Lambda(X_i, X_i', \Lambda)|X_i]$ need not be the actual expectation of the households because any function of the first period variables is a valid instrument (see Hansen, 1982; Hansen and Singleton, 1982). Here I will use forecasts from the econometrician's point of view, i.e. the fitted values from regressions on first-period variables. Overidentifying exclusion restrictions can also be constructed by using nonlinearities in the expected gradients as instruments, for example the square of the forecasted gradient.

4.5 Discussion of Identification

The theoretical restrictions in equations (21) and (22) describe how the optimal mixture of parental inputs relates to the relative prices of inputs. If the estimation was carried out in levels the estimated parameters would depend on how correlated the input mixtures are with relative prices: if large cross-sectional differences in relative prices were associated with small differences in input ratios this would be interpreted as the result of complementarity between the inputs. However, another reason for a weak relationship might be unobserved heterogeneity that is correlated with both wages and parental inputs. A key implication of the first-differences approach is that such permanent heterogeneity is eliminated and we can be much more confident that the measured relationship reflects the degree of input complementarity. The identifying variation is at the within-household level, thus it is the responsiveness of household's input ratios to changes in the relative prices they face that identifies the elasticities of substitution.

4.6 Results

Table 4 presents the results under four variations of the estimator. The main results are those presented in the first specification, but estimates do not vary greatly across columns. Standard errors are bootstrapped based on a procedure that preserves the covariance structure of the error terms. There are two sources of correlation in the error terms that need to be accounted for: (1) the cross-equation error terms of an individual may be correlated, and (2) the within equation error terms of siblings may be correlated. By resampling at the sibling group level both of these correlations are maintained.

The first important result regards the degree for substitutability between time and goods inputs, which depends on λ_0. With an elasticity of substitution just over 2.5 families can substitute purchased goods or services for their own time inputs; however, shifting the input ratio towards expenditures without changing total resources invested will reduce the net benefit to the child and skill development will suffer. Similarly, there are costs associated with reducing expenditure inputs and instead providing time inputs in the event of negative shock, e.g. an unemployment shock.

Next, the substitutability parameter λ_1 between the time inputs of spouses is only marginally less than unity. However, we would easily reject that this parameter is in fact equal to unity based on the standard errors. The quantitative importance of this parameter for single-parent households stems from its implications for returns to scale in the number of parents. At the estimated parameter value a 15% gain in effective parental time is experienced if time inputs are split equally between two spouses rather than provided by a single-parent.

The other significant parameter reflects the growth rate of the weight on expenditure inputs α_j. The level of α_j is not identified here, but the results clearly indicate that the importance of expenditure inputs increases significantly as a child gets older. We accept the hypothesis that intra-household ratios of marginal utilities have not changed in common way, i.e. we conclude that $c = 0$. Lastly, although signs of the loadings on first-differenced inverse Mill's ratios indicate positive selection, we clearly accept the hypothesis that these loadings are zero.

Finally, in comparing results across specifications we find that (1) ignoring measurement error in the inputs would bias the estimated substitution parameters upwards, and (3) including non-linearities in the expected gradients as overidentifying instruments has little effect on the point estimates.

4.7 Calibration of the Remaining Parameters

To specify the conditional probability of marriage and the distribution of fertility timing preferences I use data from the main NLSY79 cohort. For the marital status probability I fit a logit model with standardized AFQT scores as the dependent variable. I find that for every one standard deviation that AFQT increases, the log odds ratio for being married when one's first child is born increases by 1.051. For fertility timing I assume two points of heterogeneity $j^* = 22$ and $j^* = 30$. In the NLSY79 data I construct an indicator for first child born before age 26 (the median value). For every one standard deviation that AFQT increases, the log odds ratio for late fertility timing increases by 0.715.

Calibration of returns to scale in parental time ϕ_t is straightforward because time shared with siblings is observed. On average approximately 40% of the time inputs provided in families with multiple children are shared between siblings. In the CDS data this appears to be independent of the child's age. Thus, to produce one hour with each of two children a parent must spend 1.6 hours of their time endowment. Calibration of returns to scale in goods is more difficult. About half of the total spending on child related inputs are on clothing, toys/books and vacations/recreation, which are the items for which there would be some economies of scale. The other large items are individualized, i.e. school tuition, so there would not be economies of scale. Based on the equivalence scale value of 0.5 per child, and the observation that half of the inputs are scalable, I set $\phi_x = 1.5$

Four remaining parameters ω, δ, ν and α_0 were specified by matching moments of the model to features of the data. The altruism weight regulates how large the financial transfers from parents to children are. In NLSY97 data transfers from parents to young adult children total $30,566 on average (see Abbott et al. (2013)). This is 1/3 of the income reported for households in the 45-54 age bracket in SCF data, which I replicate by setting $\omega = 0.48$. The parameter δ effectively acts as the price of investing in children, and thus regulates the overall size of investments. I replicate the overall average time input received by children by setting $\delta = 23.4$. The ratio of expenditure inputs to time inputs depends on α_k. Given that time inputs are fixed by setting δ, I am able to use α_0 to replicate that parents in the CDS data spend 6% of their incomes on observed investments of goods and services. The value that achieves this is $\alpha_0 = 0.19$. Lastly, setting the parameter $\nu = 0.54$ attains an average labor supply that is 35% of the time endowment.

5 Quantitative Analysis

While the estimates of the investment function are instructive, analysis of counterfactuals scenarios is necessary in order to fully understand their implications. Two useful metrics to quantify the disadvantages faced by children from single-parent families are the average log-wage gap and the consumption equivalent welfare gap. The average log-wage gap is simply the average log-wage among individuals who grew up in traditional households minus the average log-wage among individuals who grew up in a single-parent household. The consumption equivalent welfare gap measures the proportional increase in the consumption of all who grew up with single-parents that is required to make the average expected lifetime utility of the two groups equal.[13] The counterfactual results reported in the following paragraphs refer to the values of these outcome measures in a new steady-state of the economy. Table 5 also documents these results.

The first question to address is how large the outcome gaps of children are in the benchmark economy. In terms of wages the gap is 37.4%. This is quite close to the actual gap of 40.1% observed in the NLSYC data, which is detailed in appendix C. The effect that growing up with a single-parent has on welfare is such that an 8.5% increase in annual consumption is required in order to make such individuals as well off as those who grew up in traditional households.

The second question to address is the importance of parental background. To study this I counterfactually make a person's martial status and fertility preferences orthogonal to their abilities. By doing so I ensure that the distribution of parental characteristics is independent of whether or not they are married. The proportions of married and late fertility households are held constant at the benchmark levels in this analysis. Parental background accounts for just under one-third of the measured wage gap, with a 24.7% difference in average wages remaining. Note that this is comparable to what is accounted for by parental background in the NLSYC data (see appendix C). The effect of eliminating the correlation between ability and marriage/timing has an even greater effect on welfare (proportionally) as it reduces the consumption equivalent welfare gap to only 3.8%. The reasons for this are two-fold. First, single-parents are now more able on average which reduces the wage gap of their children. Second, the lower average ability of children from single-parent households no longer disadvantages them in the marital status draw.

Next, I want to strip away the advantages that exist because of equivalence scales. For this analysis I counterfactually set the equivalence scales to $o(j, \mu, \chi) = 1$ for every household, and also eliminate scale effects in parenting by setting $\phi_x = \phi_t = 2$. The associated resource savings are a very substantial contributor to the observed wage gap, with the gap being only 15.3% after differences in equivalence scales and parental background are both accounted for. A large part of the welfare gap is also accounted for as only a 2.4% increase in annual consumption is required to equalize ex-ante welfare across parental types.

Now the really interesting experiments begin. From the analysis so far we have determined that substantial gaps in outcomes still exist even after taking out the effects of parental background and economies of scale that create the wealth differences observed across household types. The remaining experiments address the importance of input complementarities and uninsurable risk in generating the residual outcome gaps.

[13] Suppose V and V^* are the average expected lifetime values of those who grew up in single-parent and traditional household, respectively. If the annual consumption of every single-parented individual increases by $q\%$ then V will increase. The welfare gap metric measures how large q must be to make $V = V^*$. Under the utility functions assumed here one can easily show that this equals $(V^*/V)^{\frac{1}{1-\sigma}} - 1$.

The first investment complementarity I counterfactually eliminate is that between mother's and father's time inputs. That is, I set $\lambda_1 = 1$ and solve for the new steady-state. This eliminates slightly more than 5% of the residual wage gap, reducing it from 15.3% to 14.5%. The effects on the consumption equivalent welfare gap are of similar magnitude, reducing it from 2.4% to 2.3%.

To assess the importance of complementarity between parental time and expenditure investments I restore λ_1 to the estimated value and then set $\lambda_0 = 1$. This has very substantial effects, reducing the log-wage gap to 8.2% and the consumption equivalent welfare gap to 1.1%. The reason expenditure-time non-substitutability is so important is that it underlies the investment advantages implied by specialization within traditional households. With perfectly substitutable inputs single-parents can scale labor supply up or down as wages change, and the associated changes in parenting time come at no cost because of seamless replacement by/of purchased inputs. However, at the estimated degree of non-substitutability there are substantial costs associated with these replacements. All told, this result indicates that the different ways single-parent and traditional households deal with the arrival of shocks accounts for almost half of the residual wage and welfare gaps.

After non-substitutability of inputs has been accounted for the only differences that remain are due to the lack of spousal consumption insurance in single-parent households. Setting both $\lambda_0 = 1$ and $\lambda_1 = 1$ results in a 7.5% wage gap and 1.0% consumption equivalent welfare gap. This result means that the additional precautionary savings by single-parents (in lieu of spousal labor supply insurance) crowds out parental investments with the aforementioned effects.

6 Policy Analysis

6.1 A Stylized Example

Here I consider a very simple version of the model that is useful for building intuition about policy implications. This simple version is an endowment economy with no labor supply, but I think of the endowment process for single parents as riskier in order to capture the absence of the spousal labor supply insurance channel. I treat ability as a scalar produced through self-investment and goods investments. Specifically, let the skill formation technology be

$$\theta_{k+1} = \exp(\eta_k)\theta_k^{\gamma} x_k^{1-\gamma}. \tag{30}$$

Under this formulation it is straightforward to solve the dynamic programming problems of households during the skill formation periods. In the absence of policy interventions, the dynamics of investments are:

$$x_k = \gamma E_{z_{k+1},\eta_k}\left[\beta \frac{u_c(c_{k+1})}{u_c(c_k)} x_{k+1}\right] \tag{31}$$

Like any Cobb-Douglas technology, the ratio of inputs depends on their relative price and productivity. Here the price is the costs of shifting resources between ages k and $k+1$, as captured by the stochastic discount factor. The relative productivity, γ, reflects the dynamic complementarity of skill investments.

Sources of Under-Investment

Let us compare the investment dynamics in equation (31) to those that would be observed in the absence of income risk, or alternatively in a 'first-best' full insurance world:

$$x_k^{FB} = \gamma \frac{1}{1+r} x_{k+1}^{FB} \tag{32}$$

Two clear reasons why parents under-invest can be identified by comparing equations (31) and (32). First, if borrowing constraints ever bind then the stochastic discount factor in equation (31) will be smaller than the deterministic discount factor in (32). This is the usual explanation for under-investment cited by, for example, Caucutt and Lochner (2012) and Cunha and Heckman (2007).

Second, even if borrowing constraints never bind, income risk will generate a negative covariance between the stochastic discount factor and expenditure investments. To be clear the dynamics in equation (31) can be re-written using the standard expectation of a product result

$$x_k = \gamma \frac{1}{1+r} E_{z_{k+1},\eta_k}\left[x_{k+1}\right] + \gamma\beta \cdot \text{Cov}\left(\frac{u_c(c_{k+1})}{u_c(c_k)}, x_{k+1}\right) \tag{33}$$

For goods investments the negative covariance between $u_c(c_{k+1})$ and x_{k+1} will drive a wedge between current investment and future ones.

The last consideration has to do with whether $x_{k+1}(z_{k+1})$ is a concave or convex function. It would be convex if the child's abilities are used as a precautionary asset, in which case the $E_{z_{k+1},\eta_k}\left[x_{k+1}\right]$ would actually inflate current investments. However, the sizeable risk associated with the η_k terms in standard parameterizations makes this unlikely: the child's human capital is risky and unsuitable for precautionary savings purposes.

Policy Considerations

We have identified in the previous arguments that with incomplete markets the slope of child human capital investments on age is too large relative to what would be chosen when income risk is fully insured. This implies that optimal policies should aim to offset these distortions by adding appropriate wedges into household's optimality conditions. For parental expenditures it seems natural to think of a subsidy S_k^x on purchases of educational goods. Under such a policy, the dynamics of parental investments become:

$$x_k = \left(\frac{1-S^x_{k+1}}{1-S^x_k}\right)\gamma E_{z_{k+1},\eta_k}[\beta\frac{u_c(c_{k+1})}{u_c(c_k)}x_{k+1}]. \tag{34}$$

If a policy maker would like to use these instruments to induce parents to choose investments that replicate the full-insurance dynamics, then they need the wedge introduced into the right-hand-side to exceed unity. This will be the case if the sequence of age-specific subsidies is decreasing, which would cause earlier expenditures to increase by more than later expenditures. This basic logic is exactly why previous work, e.g. Cunha (2013), has found that subsidies for early subsidies are more beneficial than late ones when markets are incomplete.

Single-Parents in the Stylized Example

Now reconsider both the sources of distortion and the associated policy implication with single-parents in mind. I argued that in an endowment economy we should think of the income process of single-parents as risker because they lack insurance through spousal labor supply. If the endowment process becomes riskier then the covariance term in equation (33) becomes bigger and there is more distortion relative to the complete markets case. Because of this the corrective wedge needs to be larger, meaning that the subsidies need to be larger. One clearly needs to keep in mind that this is a partial equilibrium result that implicitly assumes no associated changes in the marriage market.

6.2 Numerical Policy Experiments

The policy experiments I present here are partial equilibrium with a balanced government budget constraint. In every experiment the tax rate of every household is increased by 0.1%, and all additional tex revenue is spent on the applicable subsidy program. This has the nice feature that the negative distortion created in each experiment is held constant.

The most important caveat associated with the partial equilibrium nature of these experiments is that marriage rates do not respond to policy changes.[14] Of course, if subsidies are targeted to single-parent families this creates an incentive for parental separation. However, we can still learn important information about the effectiveness of policy by studying the marginal effect that targeted policies have on outcomes of interest. If there are substantial partial equilibrium benefits associated with subsidizing investments in children from single-parent households then the next step is to develop a tractable model that includes endogenous marriage and fertility.[15] If the partial equilibrium benefits of programs targeted to single-parent families are small then such research would be unfounded. The

[14]Examples of how single-parenting rates might respond to policy can be found in Guner and Knowles (2009) or Bitler et al. (2004).

[15]Examples of models that include these features, but not the complex dynamics of skill formation, is Caucutt et al. (2002). One would also want explicitly consider the incentive-compatability constraint implicit in such a policy design study, as is done in, for example, related work by Gelber and Weinzierl (2012)

point is that studying the gradient of social welfare with respect to targeted policies is beneficial even if it does not provide information about general equilibrium policy effects.

Table 6 presents the results of a number of variations on the policy experiment. In all of them the excess tax revenue is spent on subsidies for expenditure investments targeted in various ways. The first experiment provides a general subsidy of 5.1% to every child independent of their age. By increasing the skill level of the economy a 0.13% gain in average labor earnings is attained, and the associated consumption equivalent welfare gain is 0.092%. Next, I experiment with a policy that is more generous for younger children, in particular one for which the subsidy rate is 50% higher for a one-year-old than a sixteen-year-old and decreases linearly. Because expenditures carry less weight in the $\mathcal{I}_k$ function when children are young, a larger average subsidy rate can be afforded. As the stylized policy example illustrated, a decreasing sequence of subsidies is beneficial. The effects on average earnings and welfare are slightly greater than for a general subsidy at 0.14% and 0.096%, respectively. Targeting to older children is less effective than a general subsidy.

The next two experiment highlight the benefits of targeting to children in single-parent households. Because they are a relatively small fraction of the population a subsidy of 14.8% to every single-parented child can be afforded. The benefits associated with this far exceed those associated with a general subsidy. The increase in average labor earnings is 0.18% and the increase in consumption equivalent welfare is 0.134%. Further refinement by targeting to younger children in single-parent households (again through a 50% decrease from one to sixteen years old) is associated with even slightly greater efficiency and welfare gains. Combined these results lead to one of the important finding of this paper: the distortionary effects that create a role for public policy affect single-parent households more, and hence policy should be more generous towards them. Again, this conclusion is subject to the marriage market general equilibrium caveat noted above.

7 Conclusion

I have developed and estimated a model of the effect that growing up in a single-parent household has on a child's adult outcomes. The primary questions of interest are whether there is an effect even when total household resources are held constant, and what the sources of any effect might be. Through counterfactual analysis I find that 40% of the gap between the outcomes of children from one- and two-parent households is due to factors above and beyond differences in resources.

The basic frictions inherent in the model are incomplete markets for insurance against economic risk and the missing market for parental time inputs. With incomplete markets the imperfect substitutability of parental time and market purchased inputs creates a disadvantage for single-parent families. When a single-parent responds to income shocks they must adjust the bundle of inputs they provide to their children much more than traditional households do. These adjustments are costly if market purchased inputs and parental time are imperfect substitutes. I estimate that the elasticity of substitution between time and goods is 2.6, and the substitution parameter is significantly less than unity (accepting the null would have implied perfect substitutes). In quantitative analysis I find that imperfect substitution of parental time and market goods accounts of nearly half of the difference in outcomes that remains after controlling for resource differences.

The structural framework also allows for the possibility that imperfect substitution of parental time inputs creates disadvantages when a parent is missing. One interpretation of this is that the higher stress levels experienced by single-parents reduces the effectiveness of their time inputs. Although I estimate that the substitutability parameter is significantly less than unity, the point estimate of the elasticity of substitution is quite large being nearly equal to six. As a result this channel plays a limited role in producing the outcome gap observed in the model. The non-substitutability of mother's and father's time inputs accounts for only 5% of the difference in outcomes that remains after controlling for resources.

The final channel to consider is the greater need for precautionary savings among single-parent households because of the absence of insurance through spousal labor supply. This channel accounts slightly less than half of the non-resource driven gap in outcomes between children from single-parent and traditional households.

Finally, I studies the partial-equilibrium effects of policy changes in both analytically, within a simplified stylized version of the model, and numerically within the full model. The fact that marriages are exogenous in the model limits interpretation of the results, but in general I find that policies targeted to children from single-parent families are more beneficial than general programs or programs targeted to younger children. The larger gradient of social welfare with respect subsidies for investments by single-parent households should motivate future research to incorporate endogenous marriage and fertility into the type of model I have developed here, and to consider deriving optimal policies subject to marital status incentive constraints.

References

Brant Abbott, Giovanni Gallipoli, Costas Meghir, and Giovanni L Violante. Education policy and intergenerational transfers in equilibrium. Technical report, National Bureau of Economic Research, 2013.

S Rao Aiyagari, Jeremy Greenwood, and Ananth Seshadri. Efficient investment in children. *Journal of Economic Theory*, 102(2):290–321, 2002.

Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The review of economic studies*, 58(2):277–297, 1991.

Marianne P Bitler, Jonah B Gelbach, Hilary W Hoynes, and Madeline Zavodny. The impact of welfare reform on marriage and divorce. *Demography*, 41(2):213–236, 2004.

Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. *The American Economic Review*, pages 1887–1921, 2008.

Richard Blundell, Luigi Pistaferri, and Itay Saporta-Eksten. Consumption inequality and family labor supply. Technical report, National Bureau of Economic Research, 2012.

Robert H Bradley and Robert F Corwyn. Socioeconomic status and child development. *Annual review of psychology*, 53(1):371–399, 2002.

Martin Browning, Pierre-Andre Chiappori, and Arthur Lewbel. Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies*, page rdt019, 2013.

Pedro Carneiro and Rita Ginja. Partial insurance and investments in children. 2012.

Pedro Carneiro, Flavio Cunha, Emanuela Galasso, and Italo López García. The role of beliefs in parental investments and child development: Evidence from a parenting intervention in chile. 2013.

Pedro Carneiro, I Lopez Garcia, Kjell G Salvanes, and Emma Tominey. Intergenerational mobility and the timing of parental income. Technical report, 2014.

Elizabeth M Caucutt and Lance Lochner. Early and late human capital investments, borrowing constraints, and the family. Technical report, National Bureau of Economic Research, 2012.

Elizabeth M Caucutt, Nezih Guner, and John Knowles. Why do women wait? matching, wage inequality, and the incentives for fertility delay. *Review of Economic Dynamics*, 5(4):815–855, 2002.

Miles Corak. Death and divorce: The long-term consequences of parental loss on adolescents. *Journal of Labor Economics*, 19(3):682–715, 2001.

Flavio Cunha. Investments in children when markets are incomplete. Technical report, 2013.

Flavio Cunha and James Heckman. The technology of skill formation. *The American Economic Review*, 97(2):31–47, 2007.

Flavio Cunha, James J Heckman, and Susanne M Schennach. Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931, 2010.

Daniela Del Boca, Christopher Flinn, and Matthew Wiswall. Transfers to households with children and child development. 2012.

Daniela Del Boca, Christopher Flinn, and Matthew Wiswall. Household choices and child development. *The Review of Economic Studies*, 81(1):137–185, 2014.

Geoffrey R Dunbar, Arthur Lewbel, and Krishna Pendakur. Children's resources in collective households: identification, estimation, and an application to child poverty in malawi. *The American Economic Review*, 103(1):438–471, 2013.

Jesus Fernandez-Villaverde and Dirk Krueger. Consumption over the life cycle: Facts from consumer expenditure survey data. *The Review of Economics and Statistics*, 89(3):552–565, 2007.

George-Levi Gayle, Limor Golan, and Mehmet A Soytas. What accounts for the racial gap in time allocation and intergenerational transmission of human capital? *unpublished*, 2013.

Alexander M Gelber and Matthew Weinzierl. *Equalizing outcomes vs. equalizing opportunities: optimal taxation when children's abilities depend on parents' resources*. National Bureau of Economic Research, 2012.

Jonathan Gruber. Is making divorce easier bad for children? the long-run implications of unilateral divorce. *Journal of Labor Economics*, 22(4):799–834, 2004.

Nezih Guner and John Knowles. Why is the rate of single-parenthood lower in canada than in the us? a dynamic equilibrium analysis of welfare policies. *Canadian Journal of Economics/Revue canadienne d'économique*, 42(1):56–89, 2009.

Nezih Guner, Remzi Kaygusuz, and Gustavo Ventura. Income taxation of us households: Facts and parametric estimates. *Review of Economic Dynamics*, 2014.

Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, pages 1029–1054, 1982.

Lars Peter Hansen and Kenneth J Singleton. Generalized instrumental variables estimation of non-linear rational expectations models. *Econometrica: Journal of the Econometric Society*, pages 1269–1286, 1982.

Jonathan Heathcote, Kjetil Storesletten, and Giovanni L Violante. The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy*, 118(4):681–722, 2010.

James Heckman. Shadow prices, market wages, and labor supply. *Econometrica*, pages 679–694, 1974.

James J Heckman and Stefano Mosso. The economics of human development and social mobility. *Annual Review of Economics*, 2014.

James J Heckman, Jora Stixrud, and Sergio Urzua. The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics*, 24(3):411–482, 2006.

Hilary W Hoynes, Diane Whitmore Schanzenbach, and Douglas Almond. Long run impacts of childhood access to the safety net. Technical report, National Bureau of Economic Research, 2012.

Greg Kaplan and Gianluca Violante. A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82:1199–1239, 2014.

Michael Keane, Robert Moffitt, and David Runkle. Real wages over the business cycle: Estimating the impact of heterogeneity with micro data. *The Journal of Political Economy*, pages 1232–1266, 1988.

John Laitner. Random earnings differences, lifetime liquidity constraints, and altruistic intergenerational transfers. *Journal of Economic Theory*, 58(2):135–170, 1992.

Michael E Lamb. *The role of the father in child development*. John Wiley & Sons, 2004.

Kevin Lang and Jay L Zagorsky. Does growing up with a parent absent really hurt? *Journal of Human Resources*, pages 253–273, 2001.

Hamish Low, Costas Meghir, and Luigi Pistaferri. Wage risk and employment risk over the life cycle. *The American economic review*, 100(4):1432–1467, 2010.

Kjetil Storesletten, Christopher I Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633, 2004.

Petra E Todd and Kenneth I Wolpin. The production of cognitive achievement in children: Home, school, and racial test score gaps. *Journal of Human capital*, 1(1):91–136, 2007.

A Co-movements of Inputs and Wages

First I illustrate the data patterns regarding the changes in human capital development inputs that are associated with changes in household circumstance from a purely descriptive view. Doing this shows how single-parent and traditional families to make quite different adjustments when circumstances change. After this I will attempt to estimate the responsiveness of inputs to unexpected wage shocks by household type, although data limitations will hinder this estimation to some degree.

The first panel of Table 7 illustrates how parental inputs vary with changes in the wages of the head of household. The main point is that the co-movements of wages and inputs are quite different between single-parent and traditional households. For traditional families the total time input a child receives does not change, but expenditure inputs increase with the heads wage. In contrast, in single-parent families parental time inputs fall significantly when the head's wage increases, and there is a much larger increase in expenditure inputs than in traditional families. One interpretation is that single-parents compensate for the lost time inputs with extra expenditure inputs. The second panel shows that within traditional families labor supply and child care time allocations adjust in such a way that the total time input received by a child is approximately constant.

Next I attempt to measure the responses of parental inputs to *unexpected* shocks to the head's wages. To do so I follow the first steps of Blundell et al. (2008) and regress logs of all variables on a set of observables: marital status, parental ages, race, family size, state, parental education, and child's age. If households condition their expectations on the same information then the residuals reflect unexpected changes in the variables. Table 8 reports the same results as Table 7, but using first difference of the residuals in place of first-differences of the actual variables.

The main findings of Table 8 are essentially the same as Table 7, but to the extent that predictable information has been purged the magnitudes of the responses have gotten smaller. The differential responses to shocks by single-parent and traditional households still appear.

Note that the full Blundell et al. (2008) estimator is not feasible here because the panel data are only across two periods, not the required three. For an example of work using the full Blundell et al. (2008) estimator on parental investment data see Carneiro and Ginja (2012).

B Labor Force Selection Equation

Specification (3) of table 4 reports the robustness of the results to the inclusion of selection terms based on inverse Mill's ratios. Keane et al. (1988) develop the result that including first-differences of the inverse Mill's Ratio in the estimating equation corrects for selection bias that results from time varying heterogeneity under the assumption that the time varying heterogeneity is *i.i.d.*. Specifically, see equation (9) of Keane et al. (1988). In this case the estimating equations are replaced by

$$\begin{aligned}\Delta\ln(\frac{1-\alpha_k}{\alpha_k})+(1-\lambda_0)\Delta\ln(\hat{x}_{k,i})+(\frac{\lambda_0}{\lambda_1}-1)\Delta\ln(\hat{t}_{1,k,i}^{\lambda_1}+\hat{t}_{2,k,i}^{\lambda_1})+ & \\ +(\lambda_1-1)\Delta\ln(\hat{t}_{1,k,i})-\Delta\ln(\tilde{w}_{1,i})-b_1\Delta\iota_{1,i} &= \epsilon_{1,i} \qquad (35)\\ (\lambda_1-1)\Delta\ln\left(\frac{\hat{t}_{2,k,i}}{\hat{t}_{1,k,i}}\right)-\Delta\ln\left(\frac{U_{c_2,i}}{U_{c_1,i}}\right)-\Delta\ln\left(\frac{\tilde{w}_{2,i}}{\tilde{w}_{1,i}}\right)-b_2\Delta\iota_{2,i} &= \epsilon_{2,i}, \qquad (36)\end{aligned}$$

where $\iota_{1,i}$ and $\iota_{2,i}$ are the inverse Mill's ratios from selection equations.

Four different probit models were fit to generate the required inverse Mill's ratios: two equations by two years. The instruments included were the state maximum unemployment benefit, which is a proxy for the generosity of the outside option, and the unearned (capital) income of the household. These are the same instruments used by Low et al. (2010). I also include a full set of child's age dummy variables. The results of these probit models are provided in Table 9. At least one of the instruments is a significant predictor of selection in each specification, and from the main results in table 4 we see that the positive nature of selection is identified, although not significantly.

C Descriptive Regressions

Using data on children from the NLSY79 child and young adult survey, and their parents from the main NLSY79 survey, I explored to associations between a child's outcomes, their parents' background and related variables. Specifically, I used a sample of 737 full time employed young adults over the age of 25 from the first wave (2012) of the young adults survey for whom complete income and parental background data was available.

Many (nearly half) of the sample of children experience some change in parental composition over the course of their childhood. For this reason I generated a fuzzy definition of overall parental status. This definition is the fraction of years prior to the age of 17 that the child lived in a two-parent household. This is equivalent to the 'number of years with an absent parent' variable used in prior research (e.g. Lang and Zagorsky, 2001). I also considered threshold rules, e.g. single-parent status assigned if more than 50% of years were spent with a single-parent, and the results were very similar. A useful aspect of the fractional approach is that the coefficient represents the effect of spending the entire childhood with two-parents versus a single-parent.[16] Parental background data consist of the mother's standardized AFQT score, average family income (mother and father) over the NLSY79 sample periods, and the age of the mother when the child was born.

Table 10 presents the results of regressions with the child's (log) personal income as the dependent variable. The first specification includes only the parental status variable and a full set of child's age dummy variables, while the remaining specifications add parental background variables. The overall

[16] This would not clearly be the case with a threshold rule approach.

income gap between children from one and two-parent households is 40.1%. More than half of this is explained by parental background variables that are correlated with household composition, particularly the mother's AFQT score.

When mother's AFQT score, total parental income and the mother's age at birth are controlled, a 19.1% wage gap remains.[17] This remaining gap is generated (at least in part) by the human capital production channels emphasized in this paper.

[17] Because these are full-time employed adults I assume their work hours are similar.

Table 1: Sample Averages by Age

Age	Sample Size	Expd.	Mom's hrs	Dad's hrs	LW-score	AP-score
5	29	1165	21.8	8.6	14.0	15.5
6	102	1965	19.8	11.6	21.7	20.5
7	104	1817	19.2	9.0	29.4	24.7
8	108	1859	19.2	8.5	36.0	28.8
9	102	2181	18.9	10.0	38.5	32.2
10	119	2191	18.0	11.0	42.8	35.3
11	183	2163	17.1	9.2	43.1	35.5
12	173	2444	16.0	9.4	44.4	37.4
13	178	2637	14.9	7.6	46.4	39.3
14	164	2769	14.5	8.9	47.0	40.7
15	143	2599	13.3	7.5	48.1	41.6
16	145	2900	12.3	6.6	48.8	42.4
17	127	3201	11.1	6.1	49.8	43.0

Notes: Data are from the child's perspective, e.g. mom's hours are hours spent with the particular child. Expenditures are annual (in real 1983 dollars), while hours are weekly. Abilities indicators are raw scores.

Table 2: Single-Parent versus Traditional

	Two-Parent HH	Single-Parent HH
Avg Child's Age	8.91	8.82
Avg Father's Time Input (hr/yr)	793	179
Avg Mother's Time Input (hr/yr)	1019	1014
Avg Expenditure Input ($/yr)	$2200	$1797
Avg Mother's Labor Supply	1142	1641
Avg Father's Labor Supply	2202	–

Notes: These data are from the 2002 wave of the CDS and exclude the 5% of single parents who are fathers. These data demonstrate the disadvantage of children in single-parent households in terms of the level of resources invested in them.

Table 3: Cross Period Variation - Estimation Sample

	Age	Expd. (1983)	Head's hrs	Wife's hrs	LW-score	N
2002	8.96	2077	856	793	36.9	543
2007	13.81	2843	676	586	47.0	543

Notes: 543 children have non-missing information for the head across both time periods. Observations with a change in the head or wife are dropped.

Table 4: Estimation Results

	(1)	(2)	(3)	(4)
Measurement Error Corrected	Y	N	Y	Y
Overidentifying Non-linearities	Y	Y	Y	N
Selection Corrections Included	N	N	Y	N
λ_0	0.6321	0.7814	0.6240	0.6200
	(0.171)	(0.073)	(0.144)	(0.134)
λ_1	0.8275	0.9064	0.8304	0.8324
	(0.079)	(0.049)	(0.078)	(0.108)
$\Delta \ln(\frac{1-\alpha_j}{\alpha_j})$	-0.1952	-0.1066	-0.2000	-0.2028
	(0.107)	(0.045)	(0.090)	(0.084)
$E[\Delta \ln(U_{c_2}/U_{c_1})]$	0.0488	0.0439	0.0519	0.0513
	(0.046)	(0.029)	(0.033)	(0.034)
b_1			0.1203	
			(0.349)	
b_2			0.4703	
			(0.697)	

Notes: Specification (1) is the benchmark. The remaining specifications present variations on this to illustrate the robustness of the estimates. The main findings are that both substitutability parameters are statistically less than unity, and there is statistically significant growth in the relative importance of expenditure as children get older. Standard errors are based on a block bootstrapping procedure that maintains the covariance structure of the error terms by resampling at the sibling group level.

Table 5: Quantitative Analysis

Experiment	Wage Gap	Welfare Gap
Benchmark	37.4%	8.5%
Uncorrelated Parental Background	24.7%	3.8%
No Economies of Scale	15.3%	2.4%
Time Inputs Perfect Substitutes	14.5%	2.3%
Market Goods and Time Perfect Substitutes	7.5%	1.0%

Notes: The wage gap is difference between average log wages of those growing up in single- versus two- parent households. The welfare gap is the consumption equivalence of the gap in the average lifetime utility of the two groups. The experiments are additive downwards as successive features are stripped away. The residual 7.5% is due to extra precautionary savings crowding out investments in children.

Table 6: Numerical Policy Analysis

Experiment	Average Subsidy Rate	Change in Average Earnings	Welfare Gain (CEV)
Constant Subsidy	5.1%	0.13%	0.092%
50% Decreasing w/age	5.6%	0.14%	0.096%
50% Increasing w/age	4.7%	0.12%	0.091%
Single Parents Constant	14.8%	0.18%	0.134%
Single Parents Decreasing	15.1%	0.19%	0.135%

Notes: In all experiments there is a level increase of 0.1% in the income tax rate. Welfare Gain (CEV) is the consumption equivalent of the change in average lifetime utility in the new steady-state. The 50% decreasing experiment involves subsidy rates that are half as large for the oldest children compared to the youngest. Policies targeted to younger children are slightly more beneficial than general policies. Policies targeted to children in single-parent households are substantially more beneficial than general subsidies. When considering policies targeted to single-parent households one needs to keep in mind that marriage is exogenous in the model.

Table 7: Co-movements of Inputs and Wages

Responses of basic inputs:

x	Two-Parents	Single-Parent
total expenditure	0.062^{**}	0.143^{**}
total parental time inputs	0.01	-0.07^{***}

Within Married Households Reallocation:

head's time input	-0.094^{**}	
wife's time input	0.088^{**}	

Notes: The co-movements of wages and inputs are quite different between single-parent and traditional households. An increase in the parent's wage rate is associated with reduced time inputs by single-parents, but a much larger expenditure input than experienced by children in two-parent families. Within traditional families labor supply and child care time allocations adjust in such a way that the total time input received by a child is approximately constant.

Table 8: Responses of Inputs of Unexpected Wage Changes

Responses of basic inputs:

x	Two-Parents	Single-Parent
total expenditure	0.045^{*}	0.133^{*}
total parental time inputs	0.00	-0.022^{**}

Within Married Households Reallocation:

head's time input	-0.083^{**}	
wife's time input	0.078^{**}	

Notes: The main findings of Table 8 are essentially the same as Table 8, but to the extent that predictable information has been purged the magnitudes of the responses have gotten smaller. The differential responses to shocks by single-parent and traditional households still appear.

Table 9: First-Stage Selection Probits

	(1)	(2)	(3)	(4)
Unearned Income	-5.2e-6	-4.2e-6	-5.9e-6	-7.9e-6
	(1.5e-6)	(2.3e-6)	(3.2e-6)	(6.7e-6)
Max UI Benefit	-7.3e-4	-6.4e-4	-3.8e-4	-1.4e-3
	(4.9e-4)	(4.9e-4)	(6.9e-4)	(6.7e-3)
Child's Age Dummies	Y	Y	Y	Y

Notes: Specifications (1) and (2) predict labor force participation of the head of household in the first and second periods, respectively. Predictions from these specifications are used to form $\Delta\iota_{1,i}$. Specifications (3) and (4) predict labor force participation of the head *and* wife in the first and second periods, respectively. Predictions from these specifications are used to form $\Delta\iota_{2,i}$.

Table 10: Descriptive Regressions

	(1)	(2)	(3)	(4)	(5)
parental status	0.401	0.211	0.310	0.346	0.191
	(0.088)	(0.101)	(0.083)	(0.088)	(0.097)
mother's AFQT		0.241			0.192
		(0.038)			(0.051)
average family income			0.171		0.124
			(0.034)		(0.040)
child's age	Y	Y	Y	Y	Y
mother's age at birth	N	N	N	Y	Y
R^2	0.055	0.102	0.094	0.093	0.159

Notes: The dependent variable is the child's log income in 2012. All specifications include a full set of child's-age dummy variables. The fourth and fifth specifications include a full set of mother's age at birth dummies.