Designing Cash Transfers in the Presence of Children's Human Capital Formation

Joseph Mullins [*][†]

June 24, 2022

Abstract

This paper finds that accounting for the human capital development of children has a quantitatively large effect on the true costs and benefits of providing cash assistance to single mothers in the United States. A dynamic model of work, welfare participation, and parental investment in children introduces a formal apparatus for calculating costs and benefits when individuals respond to incentives. The model provides a tractable outcome equation in which a policy's effect on child skills can be understood through its impact on two economic resources in the household – time and money – and the share of each resource as factors in the production of skills. These key causal parameters are cleanly identified by policy variation through the 1990s. The model also admits simple and interpretable formulae for optimal nonlinear transfers in the style of Mirrlees (1971), with novel features arising when child skill formation is accounted for. Using a broadly conservative empirical strategy, estimates imply that optimal transfers are about 20% more generous than the US benchmark, and shaped very differently. In contrast to current policies, the optimal policy discourages labor supply at the bottom of the income distribution due to the costly estimated impacts of work on child development. The finding underscores the importance of reconciling results in the literature on the developmental effects of maternal employment. Finally, a counterfactual model exercise suggests that changes to the welfare and tax environment after 1996 had negative average effects both on maternal welfare and child skill outcomes, with a significant degree of redistribution across latent dimensions.

1 Introduction

A proper accounting of the costs and benefits of cash assistance programs for households with children must include the long-run impacts on children's human capital. This paper precisely articulates and quantitatively validates the argument by estimating a model in which single mothers work, participate in government assistance programs, and shape their children's human capital through time and money investments. Accounting for the role played by time and money as economic resources has implications for the optimal size and shape of transfers to these households, which a formal analysis of optimal taxation in the spirit of Mirrlees (1971) will demonstrate.

A growing empirical literature establishes the motivating facts. First, maternal time and household income have both been shown to play a causal role in shaping child skill outcomes (Duncan, Morris, and Rodrigues, 2011; Dahl and Lochner, 2012; Akee, Copeland, Costello, and Simeonova, 2018; Bernal and Keane, 2010, 2011). Second, an individual's skills shape their life-cycle outcomes across multiple economic and social dimensions (Cunha, Heckman, and Schennach, 2010; Heckman, Stixrud, and Urzua, 2006; Heckman, Pinto, and Savelyev, 2013). Third, interventions that boost skill outcomes in childhood have been shown to reap large long-run returns (Heckman, Hyeok, Pinto, Peter, Moon, Savelyev, and Yavitz, 2010; García, Heckman, Leaf, and Prados, 2020; Bailey, Sun, and Timpe, 2021; Kline and Walters, 2016; Chetty, Friedman, Hilger, Saez, Schanzenbach, and Yagan, 2011), including those that provide material economic support to disadvantaged households (Barr, Eggleston, and Smith, 2022; Bastian and Michelmore, 2018; Hoynes, Schanzenbach, and Almond, 2016; Aizer, Eli, Ferrie, and Lleras-Muney, 2016; Bailey, Hoynes, Rossin-Slater, and Walker, 2020). Put together, existing evidence suggests that children's human capital development can shape the weighing of benefits and costs of social programs. Despite this evidence, Aizer, Hoynes, and Lleras-Muney (2022) point out that the United States Congressional Budget Office does not account for such impacts when evaluating the fiscal effects of proposed legislation; nor have economists provided a coherent framework for doing so. This paper provides such a framework for a specific class of policies (cash assistance that conditions on work, earnings, and participation behavior) for a specific population (single mothers) and uses the framework for two quantitative exercises. An optimal policy calculation à la Mirrlees (1971) offers precise insights into how accounting for children's skill formation can have quantitatively quite important implications for the social desirability of different cash assistance programs. A second counterfactual exercise evaluates the long-run skill and welfare impact of historically significant changes to US safety net policies in the 1990s.

The framework requires four key ingredients: (1) an economic model in which choices are articulated with respect to the class of policies considered; (2) a technology of skill formation that maps these choices to child human capital outcomes; (3) a mapping from skills to economic resources in the long-run; and (4) credible sources of identification for each of these three mechanisms.

Describing ingredients (1)-(3) along with their sources of empirical discipline ("identification") will provide a useful overview of this paper's methodological choices and limitations before highlighting some quantitative results.

The paper develops a dynamic model in Section 2 that carefully replicates the complex cash assistance policies available to single mothers in the United States. These fall into three categories:

1. Food stamps: the Supplemental Nutritional Assistance Program (SNAP);

2. Cash welfare: Assistance to Families with Dependent Children (AFDC) before 1996, Temporary Assistance of Needy Families (TANF) thereafter; and

3. Taxes: the Earned Income Tax Credit (EITC) and the Child Tax Credit (CTC) both specifically provide refundable and non-refundable payments to households with children.

In the model, single mothers decide whether to work and whether to participate in assistance programs. Their earnings combine with prevailing tax and benefit formulae to determine net income. They then decide how much of their net income and non-work time to invest in each of their children, which determines future cognitive and behavioral skills through a human capital production function. Thus, their decisions are dynamic: they must weigh the benefits and costs of their decisions today against the future human capital of their children. The setup follows a conceptual tradition of thinking about child development as the outcome of a human capital investment problem (Becker and Lewis, 1973; Becker and Tomes, 1976). In the model, social assistance policies shape net income and hours at home through their generosity and incentive structure, and the effect on these economic resources has a spillover to children's development through mothers' investment decisions. A key assumption in the model is that mothers are unobservably different from each other in their preferences and labor market productivity. These differences in "type" turn out to be crucial in determining winners and losers from different policy regimes, which the final exercise in the paper will demonstrate.

The Panel Study of Income Dynamics (PSID) and its Child Development Supplement (CDS) provide the necessary data (Section 3) for estimating the model (Section 4). The PSID provides longitudinal data on the work, earnings, and program participation for mothers of children in the CDS, while the CDS itself provides measures on skill outcomes and maternal time investment for children in the household. The mothers in the selected sample worked and raised their children through the 1990s, which was a period of historically significant changes in US safety net policies. Most significantly, the Personal Responsibility and Work Opportunity Reconciliation Act (PRWORA) in 1996 drastically altered the nature of cash welfare: it introduced federally mandated time limits on welfare participation, imposed work requirements on participants, and gave States legislative freedom to re-allocate funds, design program features, and change benefit formulae. The 1990s also saw several large expansions in the EITC, which significantly enhanced the financial returns to work for single mothers. A number of papers using a variety of methods have

found that these reforms played a substantial role in reducing welfare participation and increasing labor force participation during this time (Hoynes, 1996; Grogger, 2002, 2003; Meyer, 2002; Chetty et al., 2013; Chan, 2013). The time period therefore provides a useful laboratory in which to study and estimate the responsiveness of single mothers to changes in work incentives. The longitudinal dimension of the data also unveils a substantial degree of heterogeneity in mothers' costs of work and participation, as well as their earnings potential, which will prove to be critical in determining widely varying distributional impacts of these historical changes to the safety net.

As an outcome of the model's solution properties, child skills can be written as a linear function of money resources (net household income) and time resources (the mother's time at home), with a residual term that is correlated with both. The coefficients on time and money are derived from the technology of skill formation and map any policy's impacts on economic resources to skill impacts. They can be credibly identified in the presence of plausibly exogenous variation in net income and mothers' work behavior, which the policy variation in the data across states and over time provides. The model structure also provides a control function approach to identification of these parameters, which produces broadly similar point estimates with greater precision. The preferred set of results – which come from a quasi-bayesian procedure that imposes weak priors on plausible parameter values – assign a statistically significant causal impact to both time and money in skill formation, albeit with more conservative magnitudes than existing estimates in the literature.

The long-run economic value of cognitive and behavioral skills is not directly estimable from the PSID-CDS data, with CDS children having reached only young adulthood. In the absence of truly credible numbers, this paper opts for conservative ones instead, focusing on the economic value of skills from earnings and criminal behavior only. The PSID-CDS data permit estimates of the effect of cognitive and behavioral skills on earnings and criminal behavior in young adulthood. External sources of evidence then provide a means to extrapolate these effects to net present values over the life-cycle. For earnings, a simple quadratic forecasting model of earnings from Current Population Survey (CPS) data is used, while numbers are borrowed from Heckman, Hyeok, Pinto, Peter, Moon, Savelyev, and Yavitz (2010) and Heckman et al. (2013) for crime. In the latter case, the chosen numbers are adjusted downward based on a comparison of the evidence in the PSID sample to that found in Heckman et al. (2013).

The model allows one to then imagine – as this paper does in Section 5.1 – the problem of a government that wishes to maximize the welfare of single mothers subject to a net present value resource constraint. At the margin the government considers the return to offering a dollar to households for a particular choice against the cost, which consists of a mechanical component (that dollar) and a behavioral component (the change in total costs when individuals respond to this change in incentives). Classically, this is where the problem ends. In this model, the government must also consider the direct impact of that dollar on the net present value of future resources through child skills, and the behavioral impacts on the child's future human capital through moth-

ers' responses in work decisions. Estimates suggest in Section 5.2 that these considerations weigh heavier than the more traditional fiscal trade-offs. The end result is that the optimal non-linear system of transfers is far more generous and shaped quite differently from the existing system in the US.

When using the sample in the year 2000 as a benchmark, average net income under the optimal policy is \$1500/month, a number that is 20% higher than under existing policies. The implications for shape are even more severe. The average welfare participant in the benchmark received \$750 per month when not working. In the optimal policy – which is universal – all individuals receive nearly \$1500 per month when not working. The payment is then reduced to almost zero for small positive amounts of earnings, providing a stark contrast to the US benchmark, where work for low income earners is strongly encouraged through the Earned Income Tax Credit and welfare work requirements. A theoretical analysis of the optimal policy problem combined with some numerical experiments demonstrates that these provocative prescriptions for optimal work incentives are driven entirely by estimates on the importance of time relative to money. The preferred estimates imply that the developmental effects of work are negative and quite costly if potential earnings are not sufficiently high. Thus, the planner wishes to discourage work for low earning mothers. When the factor share of time is set to zero the optimal policy once again exhibits work subsidies, as in the case without children (Diamond, 1980; Saez, 2002), albeit with greater generosity due to the effect of income on future skill outcomes.

The paper concludes by studying a counterfactual in the estimated model that undoes welfare reform by pausing the policy environment in 1996. Under the counterfactual scenario, average welfare and child skills both improve modestly, suggesting that welfare reform had negative average impacts on child skills and maternal welfare. Arguably more striking are the differences in welfare and skill impacts across latent types. Welfare reform appears to have offered relatively small welfare gains to a majority of single mothers while imposing a large cost on others. The finding largely disappears when ex-ante heterogeneity in the model is under-accounted for, which emphasizes the importance of properly estimating ex-ante heterogeneity when decomposing welfare effects into insurance and redistribution components (Floden, 2001).

The quantitative exercises in this paper connect four literatures. The first seeks to discipline theories of optimal taxation with microeconomic evidence, either taking as given a redistributive objective (Saez, 2001, 2002; Blundell and Shephard, 2011) or building in a role for transfers through incomplete insurance markets (Heathcote, Storesletten, and Violante, 2017). To the extent that papers in this literature consider transfers aimed at households with children (Guner, Kaygusuz, and Ventura, 2020; Bruins, 2019; Ho and Pavoni, 2020), they do not consider the human capital spillovers to children and how this affects policy conclusions.[1] Another literature studies the

---

[1]A close counterpart that does not study optimal policy is Caucutt and Lochner (2020). The authors estimate a model of household investment in children with lifecycle borrowing constraints. They use the model to evaluate policies that

investment behavior of parents and the resultant human capital outcomes of children through a technology of skill formation, either modeling explicitly the demand for inputs as a maximization problem (Del Boca, Flinn, and Wiswall, 2014; Bernal, 2008; Griffen, 2019; Brilli, 2022; Caucutt, Lochner, Mullins, and Park, 2020; Gayle, Golan, and Soytas, 2015), or estimating directly the demand for inputs in reduced form (Cunha, Heckman, and Schennach, 2010; Attanasio, Meghir, and Nix, 2020). A third empirical literature documents direct empirical evidence on the role played by economic resources in shaping short and long-run child outcomes.[2] Finally, by using the model to study the behavior of single mothers under welfare reform, this paper repeats past efforts to understand program participation and labor supply behavior and quantify the welfare effects of policy reforms (Hoynes, 1996; Chan, 2013; Keane and Wolpin, 2010).

2   Model

The model positions mothers as the sole decision-makers in a family unit. It features a particular combination of Cobb-Douglas production and log preferences that greatly simplifies the interaction between child skills in any period and the rest of the dynamic problem, as in Del Boca, Flinn, and Wiswall (2014). The model extends the framework of Del Boca et al. (2014) in that it adopts a multidimensional skill technology, and embeds the simplification inside a dynamic, complex, policy environment. Unlike Del Boca et al. (2014), the model does not feature multiple categories of time use and considers only one parent households. Appendix C provides a general solution to the model with multiple public and private time use categories to show that the assumption of a single private time investment is inconsequential for the policies and counterfactuals considered here.

2.1   Environment and Demographics

This is a model of a sample of M mothers indexed by $m = 1, ..., M$. Time is discrete and indexed by t where $t = 0$ indicates the beginning of the decision problem for each mother and T_m indicates the last period in which a relevant decision is made by mother m (i.e. the finite horizon). Each period in the model corresponds to one calendar year in the data, although some exposition that follows in the paper scales numbers to weekly or monthly terms for ease of interpretation. Let $y_{m,t}$ indicate the calendar year correponding to the t-th period for mother m and let $t_{m,y}$ indicate the period that corresponds to the calendar year y for mother m. Each mother has a deterministic sequence of fertility outcomes in which F_m children are born, with each child indexed in chronological order of birth by $f = 1, ..., F$. Let $b_{m,f}$ be the birth period of child f for mother m, implying that their age in period t is $a(m, f, t) = t - b_{m,f}$, and let $B_m = \{b_{m,f}\}_{f=1}^{F_m}$ be the sequence of birth years for mother m.

---

alleviate borrowing constraints at different life-cycle stages.

[2]In addition to the papers cited above, Aizer et al. (2022) provide a useful review of this literature.

Children are characterized at time t by a dynamic 2-dimensional vector of characteristics, $\theta_{m,f,t} = [\theta_{m,f,t,C}, \theta_{m,f,t,B}]$, representing their stock of *cognitive* (C) and *behavioral* (B) skills. When a child reaches age 18 their skills are no longer malleable, indicating the end of the investment problem. Likewise, child f's skills are not malleable until they are born ($a(m,f,t) \geq 0$). As a short hand, let $\mathcal{A}(m,t) = \{f : 0 \leq a(m,f,t) < 18\}$ be the set of children of mother m whose skills are actively developing at time t. The number of current children younger than 18, $\tilde{F}_{m,t} = |\mathcal{A}(m,t)|$, is a relevant variable in determining taxes and welfare payments. Inclusive of these birth histories is an overall environment for mother m at time t denoted by $S_{m,t}$. The vector $S_{m,t}$ contains all variables required to summarize the mother's decision problem, such as variables that determine her market wage and variables that define the tax and transfer environment she faces. The next three subsections describe all of the components of this vector.

2.2   Endowments and Choices

At every t, each mother is endowed with 112 hours per week to spend in activities of their choosing. First, mothers choose among D discrete options, indexed by $d = 1, ..., D$ where each d corresponds to an hours choice, $H_d \in \mathcal{H}$, and a program participation choice $P_d \in \mathcal{P}$. The baseline model assumes that all working individuals work for 30 hours per week ($\mathcal{H} = \{0, 30\}$) and that participation is a binary decision ($\mathcal{P} = \{0, 1\}$), resulting in $D = 4$ total discrete choices.

In addition to the discrete choice, each mother m decides on private consumption (c_{mt}), leisure (l_{mt}), time investment in each developing child ($\tau_{m,t} = \{\tau_{m,f,t}\}_{f \in \mathcal{A}(m,t)}$) and money investment in each developing child ($x_{m,t} = \{x_{m,f,t}\}_{f \in \mathcal{A}(m,t)}$).

2.3   Technology and Government Policies

2.3.1   Wages and Earnings

Each mother faces a wage $W_{m,t}$ in the labor market that generates earnings $E_{m,t} = W_{m,t}H_{m,t}$ where $H_{m,t} = H_{d_{m,t}}$ is the hours choice associated with choice $d_{m,t}$. Wages are determined by the equation:

$$\log(W_{m,t}) = \gamma_{W,0,m} + \gamma_{W,1,m}Age_{m,t} + \varepsilon_{m,t}$$

where $Age_{m,t}$ is the age of the mother at time t and $\varepsilon_{m,t}$ is a discrete random variable that takes up to $\mathcal{E}$ values and evolves according to a first order markov process with transition matrix Π_W. Initial values of $\varepsilon_{m,0}$ are drawn from a distribution Π_W^o. The heterogeneous terms $(\gamma_{W,0,m}, \gamma_{W,1,m})$ determine fixed differences and growths in labor market productivity across mothers.[3]

---

[3]Minimum wages are not imposed for this analysis.

2.3.2 Government Policies

Let $Z_{m,t}$ be the vector of government policy rules faced by mother m at time t. This section offers a general summary of the components of $Z_{m,t}$ that determine the *net transfer* from the government, which is given by the sum of food stamps (SNAP), cash welfare (AFDC/TANF), and taxes.

Time Limits The introduction of this paper identified time limits as an important component of the changes to welfare policies introduced by PRWORA. This is modeled as a lifetime limit on cash transfers from AFDC/TANF, denoted by $\Omega_{m,t}$. Setting $\Omega_{m,t} = \infty$ indicates that time limits do not apply. When $\Omega_{m,t} < \infty$, the endogenous state variable $\omega_{m,t}$ tracks mother m's accumulated periods of welfare usage. It evolves according to:

$$\omega_{m,t+1} = \omega_{m,t} + P_{m,t}\mathbf{1}\{\Omega_{m,t} < \infty\} \tag{1}$$

with $\omega_{m,0} = 0$.

Work Requirements Another prominent feature of welfare reform was the introduction of work requirements, which mandated that all non-exempt welfare participants either meet a minimum weekly hours requirement of 30 hours, or participate in state administered job search or employment training activities. Given the inexact nature of how work requirements were enforced, this paper follows Mullins (2020), which models work requirements as an additional utility cost for participating in welfare programs when not working. This assumption is based on the lack of evidence that employment or training programs had any effect on the wages of participants, but does appear to have an effect on the number of individuals who work while participating in welfare. $R_{m,t}$ is a random binary variable that indicates whether the policy environment for mother m at time t involves work requirements. $R_{m,t}$ will feature in mother's utility (described below) but not explicitly in determining welfare payments.

Transfers The net transfer from the government to a household with participation choice P and earnings E is given by the function:

$$\mathcal{T}(S_{m,t}, P, E) = P\left(T^F(Z_{F,m,t}, E, \tilde{F}_{m,t}) + \mathbf{1}\{\omega_{m,t} < \Omega_{m,t}\}T^A(Z_{A,m,t}, E, \tilde{F}_{m,t})\right) + T^T(Z_{T,m,t}, E, \tilde{F}_{m,t})$$

where $\tilde{F}_{m,t}$ is the total number of dependent children, and the triple $(Z_{F,m,t}, Z_{A,m,t}, Z_{T,m,t})$ contains all the parameters that determine net transfers from food stamps (F), welfare (A), and income taxes (T). These vectors include variables determining gross and net income eligibility requirements, benefit standards that determine overall generosity of programs based on family size, earnings disregards for welfare payments, marginal tax rates, and child-specific tax credits such as the Earned Income Tax Credit (EITC) and Child Tax Credit (CTC). Appendix B gives a thorough description of how these functions are calculated. It is convenient to simplify the expression for

Figure 1: Transfer Policies and Tax Rates: Example

This figure shows a comparison of transfers to single mothers with two children in Florida and California in the years 1990 and 2000. All dollar values are reported in year 2000 USD. The right panel shows the participation tax rate (PTR) for each combination of state and year, while the left panel shows the monthly transfer.

net household income to the relationship $Y_{m,t} = \mathbf{Y}(d_{m,t}, S_{m,t})$ where

$$\mathbf{Y}(d, S_{m,t}) = H_d W_{m,t} + \mathcal{T}(S_{m,t}, P_d, H_d W_{m,t}).$$

Figure 1 provides a sense of the variation in incentives created by differences in policy over time and across states. It plots the monthly transfer for a single mother of two in two states – California and Florida – as a function of monthly earnings in the years 1990 and 2000. To capture incentives at the extensive margin, Figure 1 also plots the corresponding participation tax rates (PTR), showing substantial variation in work incentives across states and over time. Of particular note is the large reduction in participation rax rates over time: a consequence of welfare reform and expansions in the EITC.

2.3.3 Resource Constraints

A budget constraint and time constraint each apply to the mother's decisions:

$$c_{mt} + \sum_{f \in \mathcal{A}(m,t)} x_{m,f,t} \leq Y_{m,t} \tag{2}$$

$$l_{mt} + \sum_{f \in \mathcal{A}(m,t)} \tau_{m,f,t} + H_d \leq 112 \tag{3}$$

Although standard, these resource constraints highlight the role for government transfer policies to shape child development outcomes through the effect they have on time and money constraints in the household. The resource constraints above illustrate the direct effect, while the behavioral

assumptions described below further entail a behavioral effect through incentives embodied in the function **Y**.

2.3.4 Child Development

Technology of Skill Formation The evolution of each child skill $j \in \{C, B\}$ is determined by a Cobb-Douglas production function, which is specified for a child of age a as:

$$\theta_{m,f,t+1,j} = \exp(\mu_{\theta,m,a,j} + \eta_{m,f,t,j})\tau_{m,f,t}^{\delta_{\tau,a,j}} x_{m,f,t}^{\delta_{x,a,j}} \theta_{m,f,t,C}^{\delta_{\theta,C,j}} \theta_{m,f,t,B}^{\delta_{\theta,B,j}}, \qquad j \in \{C, B\} \tag{4}$$

where $\delta_{\tau,a,j}$ is the Cobb-Douglas share of time investment at age a for the production of skill j, $\delta_{x,a,j}$ is the equivalent for money investment, and $\delta_{\theta,j',j}$ is the share of skill j' in the production of skill j. Each mother may differ in their innate ability for raising children, represented by the time invariant productivity parameter $\mu_{\theta,m,a} = [\mu_{\theta,m,a,C}, \ \mu_{\theta,m,a,B}]$, while the vector $\eta_{m,f,t} = [\eta_{m,f,t,C}, \ \eta_{m,f,t,B}]$ is a mean zero[4] shock to total factor productivity that is independent over time and across individuals. Importantly, no such restriction is placed on $\mu_{\theta,m}$ which may be freely correlated with other heterogeneous features of the problem such as mother m's preferences and productivity in the labor market.

Initial Conditions Each child f begins with a vector of skills in their year of birth, $\theta_{m,f,b(m,f)}$, which is determined by

$$\theta_{m,f,b(m,f)} = \exp(\mu_{\theta,m,-1} + \eta_{m,f,b(m,f)-1}) \tag{5}$$

2.4 Preferences

Mothers derive utility each period from their private consumption (c), leisure (l), the skills of all children currently borne (θ), and the discrete choice (d). Outcomes are ranked by the expected discounted present value of utility and a terminal payoff that depends on the final skills of all children:

$$V_{m,t} = \mathbb{E}_t \left\{ \sum_{s=t}^{T_m} \beta^{s-t} U_{m,t}(c_s, l_s, \theta_s, d_s, \epsilon_s) + \beta^{T_m+1-t} \overline{V}_m(\theta_{T_m}) \right\} \tag{6}$$

where $\mathbb{E}_t$ is the expectation operator conditional on the mother's information set at the beginning of time t. The functional forms for these two terms are:

$$U_{m,t}(c, l, d, \theta, \epsilon) = \alpha_C \log(c) + \alpha_l \log(l) + \alpha_{\theta,m} \sum_{f:a(m,f,t)\geq 0} \sum_{j=C,B} \alpha_{\theta,j} \log(\theta_{m,f,t,j}) + \alpha_{m,t,d} + \epsilon_{m,t,d}$$

and

$$\overline{V}_m(\theta) = (1-\beta)^{-1}\alpha_{\theta,m} \sum_{f=1}^{F_m} \sum_{j=C,B} \alpha_{\theta,j} \log(\theta_{m,f,T_m,j}).$$

---

[4]This is a normalization without loss of generality.

The innate utility derived from each discrete choice d consists of two stochastic components: a persistent component $\alpha_{m,t,d}$ and an idiosyncratic component $\epsilon_{m,t,d}$. Furthermore, the vector of innate utilities from discrete choices, $\{\alpha_{m,t,d}\}_{d=1}^{D}$, is restricted in the model to take the form:

$$\alpha_{m,t,d} = -\alpha_{A,m}P_d - \alpha_{H,m}\mathbf{1}\{H_d > 0\} - R_{m,t}P_d\mathbf{1}\{H_d < 30\}\alpha_R$$

where $\alpha_{A,m}$ is a heteregeneous cost of participating in AFDC/TANF, $\alpha_{H,m}$ is a heterogeneous cost of participating in the labor market, and α_R is a global parameter that dictates the non-pecuniary cost of choosing not to meet a work requirement when participating in AFDC/TANF. This cost only applies when work requirements are in effect.

Mothers differ in terms of how they value the skills of their children and each discrete choice. In the latter case, the utility derived from each choice d consists of the two time invariant components $(\alpha_{A,m}, \alpha_{H,m})$ that are fixed for each mother over time, and a component $\epsilon_{m,t,d}$ that is identically and independently distributed over time and across mothers. The vector $\epsilon_{m,t}$ is distributed as a nested logit with the outer nest containing the program participation decision with scale parameter of 1, and the inner nest containing the work decision with scale parameter σ_H.

The heterogeneous weight on child skills $(\alpha_{\theta,m})$ is a scalar variable that scales a homogenous aggregate given by the pair $(\alpha_{\theta,C}, \alpha_{\theta,B})$. Later sections of the paper discuss identification of the model and will demonstrate that these weighting parameters need not be identified for the purposes of policy and prediction.

2.5 Information

Uncertainty about future wage realizations and utility shocks are the two key sources of risk for agents in the model. Formally, while the pair $(\epsilon_{m,t}, \varepsilon_{m,t})$ is observable at time t before decisions are made, future realizations of these variables are unknown. The same information structure is imposed for realizations of the vector of TFP shocks in child development, $\eta_{m,t}$. Section 2.8 demonstrates that these shocks enter additively in mother's utility, having no affect on decision-making. In this sense, assumptions on whether these shocks are known or forecastable have no meaningful effect on the model's solution properties.

Mothers have perfect information about the path of fertility outcomes (the number and timing of all births, B_m) as well as the path of policy variables, $\{Z_{m,t}\}_{t=0}^{T_m}$. Section 2.9.2 discusses the computational complexity of the model and how this assumption swaps the very high dimensional space spanned by $Z_{m,t}$ and B_m for the (much smaller) size of the sample, M. This simplification is necessary for model tractability, though it does have some empirical content which is discussed in Section 2.9.2.

2.6 Adult Outcomes

Each child's cognitive and behavioral skills at maturity make them more productive in the labor market and less prone to socially costly behavior. The function $\mathcal{Y}(\theta)$ measures the expected net present value in units of output of these adult outcomes and takes the form:

$$\mathcal{Y}(\theta) = \gamma_{\mathcal{Y},0} + \gamma_{\mathcal{Y},C} \log(\theta_C) + \gamma_{\mathcal{Y},B} \log(\theta_B).$$

While the domains over which cognitive and behavioral skills may bear fruit are potentially numerous, the focus of this paper will be on the private and social returns through earnings and criminal behavior. Section 4.1 presents this paper's strategy for choosing empirically reasonable (but ultimately conservative) values for $\gamma_{\mathcal{Y},C}$ and $\gamma_{\mathcal{Y},B}$, based on available data in the PSID-CDS sample combined with auxiliary life-cycle evidence in Heckman et al. (2013). Statistically modeling the net present value of skills in adulthood has two benefits. First, Section 5.1 formulates a fiscal policy problem where a social planner must internalize the future costs and benefits of child development spillovers when designing taxes and transfers. The function $\mathcal{Y}(\theta)$ is a necessary input to this problem. Second, the function $\mathcal{Y}(\theta)$ provides an interpretable scale for the latent factors (θ_C, θ_B) with which to present and discuss estimates of production parameters, as well a natural aggregation of skill impacts into dollar values when evaluating counterfactual policies.

2.7 Measurement Error

This section briefly outlines a number of measurement assumptions that are required in order to interpret available data as being generated by the economic model.

2.7.1 Measurement of Child Skills

It is well known that accounting for measurement error in skills is important for estimating the technology of skill formation (Cunha et al., 2010). Accordingly, assume that log skills $\log(\theta_j)$ are noisily measured in a linear system with two measures per skill:

$$\tilde{\theta}_{m,f,l,t,j} = \mu_{j,l} + \lambda_{j,l} \log(\theta_{m,f,t,j}) + \xi_{\theta,m,f,t,j,l}, \ j \in \{C,B\}, \ l = \{1,2\}.$$

where each error term $\xi_{\theta,m,f,t,j,l}$ is independent of all other observations.

2.7.2 Measurement of Wages and Time Investment

Wages and time investment are measured with additive error that is independent across observations. The specification for each is:

$$\log(W_{m,t}^o) = \log(W_{m,t}) + \zeta_{m,t}^W$$

and

$$\log(\tau_{m,t}^o) = \log(\tau_{m,t}) + \zeta_{m,t}^\tau$$

12

where $\zeta_{m,t}^W$ and $\zeta_{m,t}^\tau$ are each independently and normally distributed with mean 0 and variances $(\sigma_W^2, \sigma_\tau^2)$.

2.8   Dynamic Problem and Solution

Mothers choose contingent plans for consumption, leisure, time and money investment, and discrete choices, to maximize the objective written in (6) subject to intratemporal constraints on time and money and intertemporal dynamics defined by time limits, wage risk, and the technology of skill formation. This section formulates the problem recursively, and shows that it can be reduced to a dynamic discrete choice model with state vector $S_{m,t}$:

$$S_{m,t} = \{Z_{m,t}^{T_m}, B_m, \omega_{m,t}, \varepsilon_{m,t}, Age_{m,t}\}.$$

In words, $S_{m,t}$ contains the full sequence of policy variables that define food stamps, welfare, and tax policy ($Z_{m,t}^{T_m}$), the full sequence of birth outcomes (B_m), accumulated periods of welfare use ($\omega_{m,t}$), the wage shock ($\varepsilon_{m,t}$), and mother's age ($Age_{m,t}$).

Integrating out taste shocks (ϵ) and TFP shocks (η) which are both independently distributed over time, the full decision problem has a recursive formulation in terms of the state vector (S_t, θ_t):

$$V_{m,t}(S_t, \theta_t) = \int \max_{d,c,l,x,\tau} \{U_{m,t}(c, l, \theta, d, \epsilon) + \beta \mathbb{E}[V_{m,t+1}(S_{t+1}, \theta_{t+1})|S_t, \theta_t, d, x, \tau]\} \, dF_\epsilon(\epsilon) dF_\eta(\eta)$$

subject to the resource constraints (2)-(3), the transition rules governed by (1) and Π_W, and the technology of skill formation governed by (4)-(5). Given the finite horizon of the problem, the recursion can be initialized at the terminal period:

$$V_{m,T_m}(S_{T_m}, \theta_{T_m}) = \int \max_{d,c,l,x,\tau} \{U_{m,t}(c, l, \theta, d, \epsilon) + \beta \overline{V}_m(\theta_{T_m})\} \, dF_\epsilon(\epsilon).$$

An important simplification exists thanks to the joint specification of technology and preferences. Intuitively, Cobb-Douglas production implies that log of skills today enter additively into log skills tomorrow, and so on. Formally, the value function is additively separable in log skills:

$$V_{m,t}(S_t, \theta_t) = \nu_{m,t}(S_t) + \alpha_{\theta,m} \sum_{f:a(m,t,f)\geq 0} \sum_{j=\{C,B\}} \alpha_{V,a(m,f,t),j} \log(\theta_{m,f,t,j}) + C_{m,t}$$

where $\nu_{m,t}$ is a component value function described below, $C_{m,t}$ is an exogenous constant term, and each $\alpha_{V,a,j}$ is an age-specific coefficient that is defined recursively according to:

$$\alpha_{V,a,j} = \alpha_{\theta,j} + \beta \sum_{j'=\{B,C\}} \delta_{\theta,j,j'} \alpha_{V,a+1,j'}, \ j \in \{C, B\}$$

with the terminal condition: $\alpha_{V,18,j} = (1 - \beta)^{-1} \alpha_{\theta,j}$. For all $a > 18$ the equality $\alpha_{V,a,j} = \alpha_{V,18,j}$ holds since the development process has concluded and mothers continue to receive the discounted present value of their child's skills at maturity. The parameters $\alpha_{V,a,j}$ are shared across mothers but are scaled by the heterogeneous term $\alpha_{\theta,m}$. Jointly, these coefficients summarize the dynamic

incentives of the investment problem, since they embody how future skills are valued relative to private consumption and leisure today. The total returns to money and time investment are captured by two aggregate terms:

$$\Gamma_{x,a} = \beta \left(\delta_{x,a,C} \alpha_{V,a+1,C} + \delta_{x,a,B} \alpha_{V,a+1,B} \right), \qquad \Gamma_{\tau,a} = \beta \left(\delta_{\tau,a,C} \alpha_{V,a+1,C} + \delta_{\tau,a,B} \alpha_{V,a+1,B} \right). \qquad (7)$$

For example, a log unit increase in time investment for a child of age a results in a skill return that has a net present value return to mother m of $\alpha_{\theta,m} \Gamma_{\tau,a}$. The constant $C_{m,t}$ is invariant to all decisions and policies and summarizes the net present value contribution of mother m's TFP in child production to future utility:

$$C_{m,t} = \alpha_{\theta,m} \beta \sum_{-1 \leq a(m,f,t) \leq 17} \sum_{j=\{C,B\}} \alpha_{V,a(m,f,t+1),j} \mu_{\theta,m,a(m,f,t),j}.$$

Finally, the component value function ν takes the form

$$\nu_{m,t}(S_t) = \int \max_d \left\{ u_{m,t}(\mathbf{Y}(d, S_t), d) + \alpha_{m,t,d} + \epsilon_d + \beta \mathbb{E}[\nu_{m,t+1}(S_{t+1}) | S_t, d] \right\} dF_\epsilon(\epsilon) \qquad (8)$$

subject to the transition rules that govern accumulated welfare use (equation (1)) and wages (transition matrix Π_W). Since F_ϵ is specified as a nested logit, the integral in (8) has an analytical solution, as does the probability of any choice d given the state S_t. The function $u_{m,t}$ is indirect utility given by the solution to the maximization problem:

$$u_{m,t}(y, d) = \max_{c,l,x,\tau} \left\{ \alpha_C \log(c) + \alpha_l \log(l) + \alpha_{\theta,m} \left(\sum_{f \in \mathcal{A}(m,t)} \Gamma_{x,a(m,f,t)} \log(x_f) + \Gamma_{\tau,a(m,f,t)} \log(\tau_f) \right) \right\}$$

subject to the resource constraints:

$$c + \sum_f x_f \leq y, \qquad l + \sum \tau_f + H_d \leq 112.$$

Since all terms are in logs, the solution to this maximization problem is a set of linear investment equations:

$$\frac{x_{m,f,t}}{y} = \left(\alpha_C + \alpha_{\theta,m} \sum_{f' \in \mathcal{A}(m,t)} \Gamma_{x,a(m,f',t)} \right)^{-1} \alpha_{\theta,m} \Gamma_{x,a(m,f,t)} \qquad (9)$$

$$\frac{\tau_{m,f,t}}{112 - H_d} = \left(\alpha_l + \alpha_{\theta,m} \sum_{f' \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f',t)} \right)^{-1} \alpha_{\theta,m} \Gamma_{\tau,a(m,f,t)} \qquad (10)$$

that results in the following expression for indirect utility:

$$\begin{aligned} u_{m,t}(y, d) &= b_{m,t} + \left(\alpha_C + \alpha_{\theta,m} \sum_{f' \in \mathcal{A}(m,t)} \Gamma_{x,a(m,f',t)} \right) \log(y) \\ &\quad + \left(\alpha_l + \alpha_{\theta,m} \sum_{f' \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f',t)} \right) \log(112 - H_d) \\ &= b_{m,t} + \tilde{\alpha}_{C,m,t} \log(y) + \tilde{\alpha}_{l,m,t} \log(112 - H_d) \qquad (11) \end{aligned}$$

14

where $b_{m,t}$ is a constant term that depends on the fractions of time at home and net income that are spent on each activity. Since this term is invariant to all decisions and environments, it can be discarded when solving the model and evaluating counterfactuals. The terms $\tilde{\alpha}_{C,m,t}$ and $\tilde{\alpha}_{l,m,t}$ are heterogeneous effective weights on consumption and leisure that vary over time depending on the relative importance of time and money investment for each child in the household and the mother's weight ($\alpha_{\theta,m}$) on child outcomes. Appendix C illustrates how in a general setup the value function at t can inherit the additive separability property from the value function at $t+1$. It also derives the solution for a more general set of public and private time use categories to show that the assumption of one investment category each for time and money satisfies Marschak's maxim (Marschak, 1953) in that only one category is sufficient for analyzing the set of policies outlined in this paper.

This section concludes by formally stating the choice probabilities given by the nested logit assumption on the taste shocks ϵ. To do so, let $\mathcal{P}(p) = \{d : P_d = p\}$ for $p \in \{0, 1\}$ be the "nests" that partition the choice set into values of d corresponding to each participation choice. Choice probabilities take the form:

$$P_{m,t}[d|S] = \frac{\left(\sum_{d' \in \mathcal{P}(P_d)} \exp(\sigma_H^{-1} v_{m,t}(d', S))\right)^{\sigma_H - 1}}{\sum_{p \in \{0,1\}} \left(\sum_{d' \in \mathcal{P}(p)} \exp(\sigma_H^{-1} v_{m,t}(d', S))\right)^{\sigma_H}} \exp(\sigma_H^{-1} v_t(d, S))$$

where $v_{m,t}(d, S)$ is the choice-specific value:

$$v_{m,t}(d, S) = u_{m,t}(\mathbf{Y}(d, S), d) + \alpha_{m,t,d} + \beta \mathbb{E}[\nu_{m,t+1}(S')|S, d].$$

2.9   Empirical Content of the Model

2.9.1   Child Outcomes

Combining the linear investment rules (9) and (10) with the production technology (4) gives a linear outcome equation for the skill vector θ that is convenient both for identification and for understanding model counterfactuals:

$$\log(\theta_{m,f,t+1}) = \mu_{\theta,m,a} + \delta_{x,a} \log(Y_{m,t}) + \delta_{\tau,a} \log(112 - H_{m,t}) + \delta_\theta \log(\theta_{m,f,t}) + e_m(a, \boldsymbol{a}) + \eta_{m,f,t} \quad (12)$$

where $a = a(m, f, t)$ is the age of child f for mother m at time t and $e_m(a, \boldsymbol{a})$ is a residual that contains the marginal propensities to invest time and money in child f given the ages of all children in the household, $\boldsymbol{a}$:

$$e_m(a, \boldsymbol{a}) = \delta_{x,a} \log\left(\left(\alpha_C + \alpha_{\theta,m} \sum_{a' \in \boldsymbol{a}} \Gamma_{x,a'}\right)^{-1} \alpha_{\theta,m} \Gamma_{x,a}\right) +$$

$$\delta_{\tau,a} \log\left(\left(\alpha_C + \alpha_{\theta,m} \sum_{a' \in \boldsymbol{a}} \Gamma_{\tau,a'}\right)^{-1} \alpha_{\theta,m} \Gamma_{\tau,a}\right). \quad (13)$$

Since the error term is a function of preferences only and invariant to policy, equation (12) demonstrates that the effects of any welfare policy reform on child outcomes can be calculated by forecasting its effects on labor supply and net household income. This result allows transparency both in how key production parameters are identified (which Section 4.4 will discuss) and how changes to welfare policy affect child outcomes in the model. Since this model focuses on labor supply at the extensive margin, the parameter δ_τ effectively determines a ceterus paribus effect of maternal employment on skills. Thus, even though the model does not feature childcare choices, it allows for an interpretation of this employment effect as the effect of maternal time relative to whatever care option is solicited when mothers work. As long as this option is invariant to the counterfactuals of interest, such an interpretation is robust.

2.9.2 Model Complexity and Dynamic Empirical Content

Two assumptions combine in this model to make it sufficiently tractable for estimation. The first is to assume that the full sequences of policy variables and birth outcomes are known. The vector of policy variables, $Z_{m,t}$, is very high-dimensional relative to what is typically computationally feasible[5], as is the space of potential birth histories, B_m. By assuming that these sequences of variables are known, the dimension of the state space for any single agent is reduced to $\sum_{t=1}^{T_m} \mathcal{E} \times (\Omega_{m,t} + 1) \times \mathbb{R}^2$, and any evaluation of an estimation criterion that uses all data points therefore requires a solution to the approximate order of $\sum_{m=1}^{M} \sum_{t=1}^{T_m} \mathcal{E} \times (\Omega_{m,t} + 1) \times \mathbb{R}^2$. The combined dimension of the state space for $Z_{m,t}$ and B_m can essentially be traded for sample size, M, which is much smaller.

The second key simplification comes from a pair of assumptions: that preferences over child skills take a log form, and child skills combine with aggregate investment in a Cobb-Douglas form to produce future skills. According to the previous section the value function is then additively separable in log-skills, further reducing the dimension of the state space for an individual problem to $\sum_{t=1}^{T_m} \mathcal{E} \times (\Omega_{m,t} + 1)$, which delivers a sufficiently tractable case when the model must be solved to the order of sample size M for the purposes of estimation.

While these assumptions are necessary for tractability, they do impose some *a priori* empirical content on the model, mainly stemming from the additive separability between child skills and other dynamic state variables. In periods prior to welfare reform when time limits do not apply, all of the dynamics of the model are loaded into the recursive coefficients on consumption and leisure ($\tilde{\alpha}_{C,m,t}, \tilde{\alpha}_{l,m,t}$), which implies that choices are unaffected by the agent's beliefs about future policy environments or births. Hence, there are no anticipatory effects from policy changes. Based on the model's fit of the data (discussed in Section 4.5.1), it does not appear that these anticipatory

---

[5]Recall that this vector includes (but is not limited to) the full set of federal and state marginal tax rates as well as income brackets, state poverty guidelines and benefit standards for up to 4 different family sizes, fixed and variable earnings disregards, and parameters determining income eligibility tests. All of these are continuous variables.

effects are necessary to fit observed time series patterns in labor supply and program participation.

Of course, anticipation effects would reappear with the introduction of additional ingredients that either augment the dynamics of the model, or break the additive separability result. The next section addresses the omission of three ingredients in particular that would change dynamics. With regards to preferences and technology, there is additional empirical content beyond the simplification of dynamics. On the developmental side, the elasticity of substitution between the stock of skills and aggregate investment is restricted to be one, which imposes a particular form of complementarity. However, previous research has shown that the Cobb-Douglas specification is difficult to reject in the data (Cunha et al., 2010; Attanasio et al., 2020). On the behavioral side, the model's implications are quite a bit stronger: it produces linear investment equations that do not depend on the child's current stock of skills. This implication is testable, and Section 4.7 uses within-family variation in skills to show that the data cannot reject it. Section 4.7 also uses data to test for substitution patterns that would violate the simple form of demand for inputs. Here again, the model cannot be rejected.

2.9.3 Other Limitations

The model's main advantages – transparent identification of key causal parameters, tractable policy analysis, and rich unobserved heterogeneity – arrive with limitations that merit discussion. Mothers cannot self-insure through savings, they do not make fertility or household formation decisions, they do not accumulate human capital of their own, and they do not make childcare arrangements. Each of these simplifying assumptions has its own motivation, but a unifying theme is that the model focuses on capturing the first order features of the data and the two choices which are most responsive to the class of policies considered: work and program participation (Meyer, 2002; Grogger, 2002).

No Savings   The ability to borrow and save would re-introduce expectations about future policies as a state variable that affects choices. Section 3 will show that individuals in this sample hold only a very small amount of cash assets, which is consistent with evidence that low income households have a marginal propensity to consume that is close to one (Patterson, 2022). Still, this caveat should be kept in mind and the upcoming policy exercises do not emphasise the effect of policy reforms on consumption insurance.

Fertility and Marriage   The choice to abstract away from fertility and household formation is largely motivated by tractability concerns: including either would likely demand simplification elsewhere in the model, as discussed in the previous section. The advantages to doing so are also unclear, since studies of fertility and marriage find either no response to welfare reform (Gennetian and Knox, 2003; Kearney, 2004) or evidence only when it comes to selection into the sample of

single mothers (Bitler et al., 2004; Low et al., 2018). Modeling this selection and the joint decision-making of single mothers with partners who are not necessarily the biological parents of all children is clearly an important challenge for future research, but beyond the current frontier.

Returns to Experience Returns to experience for mothers would add dynamics to the model and certainly modify policy conclusions regarding optimal work incentives. Section 4.7 looks for evidence of returns to work experience in the data and finds no evidence of returns for this population, echoing related results (Blundell et al., 2016).

Childcare If childcare expenditures are complementary with work in household utility, it opens up a role outside of Atkinson and Stiglitz's (1976) classic theorem for childcare subsidies as a redistributive fiscal tool, as in Ho and Pavoni (2020). When considering the human capital development of children it would be important to model not just the cost but also the quality and type of care solicited (centre-based vs in-home vs informal) when studying child-care subsidy policies.[6] The study of cash transfer and child-care policies jointly is an important avenue for future research and would require better data on childcare quality as well as additional sources of variation to credibly identify the role this channel plays in skill formation. The upcoming analysis of optimal policy identifies the total effect of employment on skills as one of two key channels through which child outcomes affect optimal policy. While equation (12) interprets this effect as being determined by the factor share of time (δ_τ), one could alternatively interpret it as the relative value of mother's time compared to the care option they use while working. Under this interpretation it is important to assume that the quality of this care is invariant to the cash assistance policies that this paper studies. To moderate concerns about this channel, Section 3 uses time diary data and finds that children in this sample overwhelmingly spend their time in informal care.

2.9.4 Latent Heterogeneity

The model allows for persistent latent heterogeneity in the costs of work ($\alpha_{H,m}$), costs of welfare participation ($\alpha_{A,m}$), the relative preference weight on child skills ($\alpha_{\theta,m}$), labor market productivity ($\gamma_{W,0,m}, \gamma_{W,1,m}$), and parenting ability ($\mu_{\theta,m}$). Allowing for heterogeneity in these dimensions precludes the use of most cross-sectional variation to estimate the model, and realistically frames the problem of credible inference. For example, one cannot simply learn the effect of welfare participation on work behavior by comparing participants to non-participants, due to self selection along unobserved dimensions. Nor can cross-sectional variation in wages be used to estimate the responsiveness of work behavior to labor supply incentives, due to latent statistical dependence between labor market productivity and costs of work.

While welfare and tax policies through the 1990s in the US will provide very useful variation

---

[6]Griffen (2019) uses high quality panel data from the ECLS-B to estimate these objects.

to estimate key parameters of the model, a panel dimension of the data is also required to uncover the rich latent heterogeneity underlying it. A *grouped heterogeneity* framework, represented by a finite mixture model, makes this problem tractable. It specifies a finite number of types, K, and an assignment for each individual $k(m) \in \{1, 2, ..., K\}$ such that:

$$(\alpha_{H,m}, \alpha_{A,m}, \alpha_{\theta,m}, \gamma_{W,0,m}, \gamma_{W,1,m}, \mu_{\theta,m}) = (\alpha_{H,k(m)}, \alpha_{A,k(m)}, \alpha_{\theta,k(m)}, \gamma_{W,0,k(m)}, \gamma_{W,1,k(m)}, \mu_{\theta,k(m)}).$$

The grouping assumption provides a flexible approach to modeling unobserved heterogeneity in the data. Section 4.2 discusses identification and estimation of the model with grouped heterogeneity. As a matter of notation, let α_A refer now to the K-dimensional vector of type-specific parameters $\alpha_{A,k}$, and similarly for all other type-specific parameters in the model.

3   Data

The *Panel Study of Income Dynamics* (PSID) and its *Child Development Supplement* is the main source of data for analysis in this paper. The PSID is a dynastic, longitudinal survey taken annually from 1968 to 1997, and biennially since 1997. The main interview provides measures of household members' earnings and program participation, as well as relationships between household members. This allows construction of a complete marriage and fertility history, in addition to an incomplete history of work and participation decisions. Measures of other demographics such as race, education, and state of residence, are also available.

The CDS consists of three waves, collected in 1997, 2002 and 2007. Any child in a PSID family between the ages of 0 and 12 at the time of the 1997 survey was considered eligible. These surveys contain a broad array of developmental scores in cognitive and socioemotional outcomes as well as information on the home environment of the child. One important feature of the survey is the availability of time use data, which is collected by the participants' completion of time diaries. The next section provides further details.

3.1   Description of Variables and Sample Selection

The CDS is comprised of several questionnaires. This paper uses two in particular: the child interview and the *primary caregiver* (PCG) interview. The Letter-Word (LW) and Applied Problems (AP) modules of the Woodcock-Johnson Aptitude test provide two measures of cognitive ability for children aged 3 and older. Two scales that measure *externalizing* (BPE) and *internalizing* (BPN) behavioral problems are used to measure the behavioral skills of children. This gives, in total, four noisy measures of child attributes that track human capital outcomes (two for each skill).

Finally, the CDS asks participant children to fill out a "time diary". This portion of the survey requires participants to record a detailed, minute by minute timeline of their activities for two days of the week: one random weekday and one random day of the weekend. Activities were

subsequently coded at a fine level of detail. When necessary, children are assisted in completion of the time diary by the PCG. These diaries provide a unique snapshot into the daily life of the child. Maternal time investment is measured in this paper by taking a weighted sum[7] of the total hours of time use in which the mother is recorded as actively participating in each diary activity.

Historically, single mothers have been the overwhelming majority of participants in child-related transfers, and are considered here as the population of interest. Since family structure is dynamic in reality, there is no perfect way to define this population in the panel dimension. This paper defines the sub-sample of single mothers to be all mothers of CDS children who were unmarried at the time of their first birth. While this model does not consider the later marriage and cohabitation decisions of single mothers, experimental studies find no evidence to suggest that these decisions are responsive to changes in welfare policies (Gennetian and Knox, 2003). At the very least, the magnitudes of the response are likely too small to be meaningfully statistically detected.

Applying this sample restriction results in a sample of 955 mothers and 1,390 children. Table 1 provides some summary statistics from this sample. Unsurprisingly, the selected sample is quite disadvantaged: mothers in this sample have low levels of education (74% have at least a High School diploma or equivalent, while 26% have less, with only 4% having obtained a Bachelor's degree) with modest earnings among those working. Additionally, we see that this sample is heavily reliant on welfare, with 48% having reported welfare use at least once, and 16% having used welfare for at least 5 years during the observed panel.

Combining the main interview of the PSID and the CDS supplement provides a panel of mothers' work and program participation decisions, earnings, fertility outcomes, time investment in children, and human capital outcomes for those children. This is sufficient to estimate the model, as Section 4 will demonstrate. However, in order to anchor child skill outcomes to a net present value in economic resources, these skills must be linked with adult outcomes. The *Transition into Adulthood Supplement* (TAS) provides follow-up survey data on CDS children once they reach adulthood. Section 4.1 below uses data from this supplement on earnings and criminal behavior. In particular, a measure of total arrests is derived by summing the TAS variable reporting the number of arrests since last interview.

The omission of savings and childcare from the model may cause concern when interpreting estimates and analyzing counterfactuals. To assess the severity of this issue, Table 1 reports total cash assets for households in the sample using PSID's Wealth Module in 1999, along with time spent in formal care according to the 1997 time diaries. The majority of sample households report no cash assets, and even individuals at the 75th percentile have modest savings. Similarly, it appears that 91% of children report spending no time in a formal care arrangement in 1997 (85% for children under the age of 5). While the scarce use of formal care may help with interpretation of estimates (one should think about the exercise as identifying the value of maternal time, holding

---

[7] $\frac{5}{7}$ for the weekday, and $\frac{2}{7}$ for the weekend

outside options fixed), there is ultimately no guarantee that these choice variables are invariant to the policies we consider here.

4 Identification and Estimation

4.1 Measuring and Anchoring Skills

Section 2.7 specifies that skills are only observable with additive error in a linear measurement system. Having two measures of each skill provides a straightforward strategy for overcoming the challenges that measurement error presents for estimating the technology of skill formation (Cunha and Heckman, 2008). Appendix D.1 describes how the parameters of the measurement system in Section 2.7 can estimated to construct two factor scores for each skill:

$$\hat{\theta}_{m,f,1,t,j} = \tilde{\theta}_{m,f,1,t,j} - \hat{\mu}_{j,1}, \qquad \hat{\theta}_{m,f,2,t,j} = \frac{\tilde{\theta}_{m,f,2,t,j} - \hat{\mu}_{j,2}}{\hat{\lambda}_{j,2}}.$$

In the limit, each factor score obeys the relationship:

$$\hat{\theta}_{m,f,l,t,j} = \log(\theta_{m,f,t,j}) + \hat{\xi}_{\theta,m,f,l,t,j} \tag{14}$$

where the measurement error terms $\hat{\xi}_\theta$ are mutually independent.

There are several strategies for giving the latent factors an interpretable scale. A common strategy is to use fractions of a standard deviation in test scores (Dahl and Lochner, 2012; Bernal and Keane, 2011; Duncan et al., 2011) while Cunha and Heckman (2008) and Cunha et al. (2010) advocate for the use of adult outcomes as a way to "anchor" the scale of unobserved skills. This paper builds in a natural anchoring strategy by positing that the net present value of skills, in terms of economic resources, can be written as:

$$\mathcal{Y}(\theta) = \gamma_{\mathcal{Y},0} + \gamma_{\mathcal{Y},C} \log(\theta_C) + \gamma_{\mathcal{Y},B} \log(\theta_B).$$

While there is not sufficient data in the PSID-CDS sample to estimate the full life-cycle returns to skills, it is possible to make reasonable (if somewhat conservative choices) by combining the available evidence on skill returns in young adulthood with external sources of evidence. This paper focuses on the economic benefits of skills on lifetime earnings and criminal behavior. In their studies of the returns to early childhood education programs, Heckman et al. (2010) and García et al. (2020) find that these two dimensions explain the majority of the programs' benefits. As with several other steps made in this section, restricting attention to these two categories also errs on the side of a conservative estimate of future returns. Formally, assume that

$$\mathcal{Y}(\theta) = \mathcal{Y}^{EARN}(\theta) - \mathcal{Y}^{CRIME}(\theta)$$

with relationships that obey:

$$\mathcal{Y}^{EARN}(\theta) = \gamma_{E,0} + \gamma_{E,C} \log(\theta_C) + \gamma_{E,B} \log(\theta_B) \tag{15}$$

$$\mathcal{Y}^{CRIME}(\theta) = \gamma_{CR,0} + \gamma_{CR,C} \log(\theta_C) + \gamma_{CR,B} \log(\theta_B) \tag{16}$$

Table 1: Descriptive Statistics

	Value
Mother	
Annual Earnings (25th percentile)	6,969
Annual Earnings (50th percentile)	15,199
Mean Annual Earnings	17,566
Cash Assets (50th percentile)	0
Cash Assets (75th percentile)	821
Cash Assets (90th percentile)	3054
Used Welfare Once (%)	47.9
Used Welfare $\geq$ 5 years (%)	15.9%
< High School (%)	26.3
$\geq$ High School (%)	73.7
Bachelor's (%)	4.35
Mean No. Children	2.86
Mean Panel Length	16.8
M	955
Child	
No Formal Care (%)	90.8
No Formal Care - Age $\leq$ 6 (%)	84.9
N	1,390

All dollar amounts are deflated to 2000 values, with income variables reported in monthly averages. When applicable, standard deviations of variables are reported in the right hand column. Cash Assets are collected at the household level from the 1999 Wealth Module, while Childcare time is measured using the 1997 time diary.

Table 2: Values (in year 2000 $USD per standard deviation) of Cognitive and Behavioral Skills

Skill	Earnings	Crime	Total
Cognitive	$\gamma_{E,C} = \$93,000$	$\gamma_{CR,C} = 0$	$\gamma_{\mathcal{Y},C} = 93,000$
Behavioral	$\gamma_{E,B} = \$47,500$	$\gamma_{CR,B} = \$9,000$	$\gamma_{\mathcal{Y},B} = \$55,500$

This table shows the chosen values for anchoring skill measures in the data according to dollars per standard deviation. Earnings values are net present values over the lifecycle, using a discount rate of 2%. They are chosen by combining estimates of the earning effects of cognitive and non-cognitive skills on the earnings of CDS children between the ages of 24 and 27 (Table 9), and using CPS data to construct a lifecycle forecast of earnings. See section D.2 for more details. The crime value of behavioral skills is chosen as one tenth of the value conservatively inferred from Heckman et al. (2013), using the estimated impact of these skills on CDS children in young adulthood as a guide (see Table 10 and section D.2.3 for more details).

To estimate the effect of skills on lifetime earnings, this paper follows Chetty et al. (2011) and Kline and Walters (2016) in assuming that the effect of skills on earnings is proportional at each age. If the percentage impact of skills on income at the beginning of the lifecycle can be estimated, this assumption allows a forecast to be constructed. As pointed out by Chetty et al. (2011), this assumption is somewhat conservative since it is likely that increases in human capital could increase the growth in earnings over the lifecycle. Analysis of the earnings of CDS children in young adulthood suggests that a standard deviation increase in cognitive skills leads to a 23% increase in earnings, while a standard deviation increase in behavioral skills leads to a 12% increase. A simple forecasting method from representative CPS samples suggests that the net present value of earnings for this cohort is $723,298, and that the CDS sample averages earnings at around 54% of this representative sample. Putting this together implies that a standard deviation increase for children in the CDS sample is worth about $93,000 in earnings for cognitive skills and $47,500 for behavioral skills.

For the case of criminal behavior, analysis of early adulthood data suggests that a standard deviation increase in behavioral skills leads to 16% fewer arrests, which is just over 10% of the magnitude of the effect size in the sample used by Heckman et al. (2013). Since that paper finds that a standard deviation in behavioral skills is worth about $90,000 in crime reduction, it is assumed that a standard deviation in behavioral skills for the CDS sample is worth about one tenth of this value ($9,000).

Appendix D.2 provides a more formal discussion of the assumptions on outcomes that lead to these calculations, in addition to technical details on the calculation of effect sizes. Table 2 reports the chosen values of the anchoring coefficients as described in this section.

4.2 Identification

Two equations from Section 2 organize the discussion of identification. Equation (11) provides a derivation of indirect utility:

$$u_{m,t}(y,d) = b_{m,t} + \left(\alpha_C + \alpha_{\theta,m} \sum_{f' \in \mathcal{A}(m,t)} \Gamma_{x,a(m,f',t)} \right) \log(y) + \left(\alpha_l + \alpha_{\theta,m} \sum_{f' \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f',t)} \right) \log(112 - H_d)$$

where $b_{m,t}$ is a policy invariant additive constant that may be ignored. Second, the model's linear investment equations produce linear dynamics in the two-dimensional vector of skills $\theta_{m,f,t}$:

$$\log(\theta_{m,f,t+1}) = \delta_{\tau,a} \log(112 - H_{m,t}) + \delta_{x,a} \log(Y_{m,t}) + \delta_\theta \log(\theta_{m,f,t}) + \mu_{k(m),\theta,a} + \eta_{m,f,t} + e_{k(m)}(a, \boldsymbol{a}) \tag{17}$$

where $a = a(m,f,t)$ is the age of the child at time t and $\boldsymbol{a} = \{a(m,f',t)\}_{f' \in \mathcal{A}(m,t)}$ is the age of all other children in the household. While the vectors Γ_x and Γ_τ are known functions of the deeper parameters $\{\delta_{x,j}, \delta_{\tau,j}, \alpha_{\theta,j}\}_{j \in \{x,\tau\}}$, it is possible to identify them directly from the data without exploiting such restrictions. This has the added benefit of not requiring that mothers have exact knowledge of the technology of skill formation to identify parameters governing their behavior.[8] It is therefore possible to consider the identification of parameters that govern agents' behavior separately from the identification of production parameters.

The presence of latent heterogeneity (indexed by k) in costs of work and program participation $(\alpha_{A,k}, \alpha_{H,k})$, labor market productivity $(\gamma_{W,0,k}, \gamma_{W,1,k})$ and investment behavior $(\alpha_{\theta,k})$, prohibits the direct use of cross-sectional variation to estimate key parameters. The data offer three separate solutions to this challenge: (1) a sufficiently long panel dimension with sufficient variation in state variables (Kasahara and Shimotsu, 2009); (2) a sufficient number of observables that "measure" latent type (Bonhomme et al., 2016); and (3) exogenous variation in policies. Appendix D.3 provides a formal discussion of how multiple identification results from the literature apply here, guaranteeing identification of the distribution of latent variables $(\Pi_K, \Pi_{W,0}, \Pi_W)$ and reduced form choice probabilities $P_{k,t}(S)$. The model imposes lower dimensional parametric restrictions on choices with straightforward relationships that can now be conditioned on type. Recall that time investment is given by

$$\frac{\tau_{m,f,t}}{112 - H_d} = \left(\alpha_l + \alpha_{\theta,k(m)} \sum_{f' \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f',t)} \right)^{-1} \alpha_{\theta,k(m)} \Gamma_{\tau,a(m,f,t)},$$

which identifies the vector Γ_τ, as well as each $\alpha_{\theta,k}$ up to the scale of α_l. Next, the type-specific parameters $(\alpha_{A,k}, \alpha_{H,k})$ are pinned down by mean rates of work and labor force participation, as are Γ_x and α_l by mean rates of work and participation for different numbers and ages of children in the household. The parameters α_C and σ_H set the scale of utilities and the correlation between preference shocks for program participation and work. They are global parameters that crucially

---

[8]Under the specification above, a scale normalization for the triple $(\alpha_\theta, \Gamma_x, \Gamma_\tau)$ is available, allowing one to set $\Gamma_{\tau,0} = 1$ without loss of generality.

determine the elasticity of work and program participation with respect to changes in policies, and are therefore identified by variation in incentives induced by these changes. The non-pecuniary cost imposed by work requirements α_R is identified by changes in the probability of working while participating in welfare induced by the introduction of work requirements. In principle, since the model exhibits two-period finite dependence once time limits have been introduced, the discount factor β is also identified (Arcidiacono and Miller, 2020) however in practice estimation of this parameter proved difficult. Accordingly, assume throughout that $\beta = 0.98$.

The derivation in Equation (17) is convenient because it presents the issue of identification in terms of the familiar linear model: one must find instruments for income and hours worked that are uncorrelated with the residual terms in this outcome equation. There are two sources of endogeneity that prohibit the use of observed net income and hours as instruments. First, the residual $e_{k(m)}(a, \boldsymbol{a})$ appears due to the fact that investments are unobserved, and is correlated mechanically with work decisions and income through the coefficients $\alpha_{\theta,k(m)}\Gamma_x$ and $\alpha_{\theta,k(m)}\Gamma_\tau$ which determine both work and investment decisions. Second, no assumption has been made to guarantee that mother's TFP, $\mu_{k(m),\theta}$, is not correlated with other heterogeneous parameters that determine behavior. It is plausible for example that mothers who find employment less costly or are more productive in the labor market may also have innately higher parenting ability.

This paper pursues two solutions. First, if each mother's true unobserved type, $k(m)$, were known, then one could instead use a flexible function of k and $(a, \boldsymbol{a})$ to control for the endogenous residual. This suggests the moment condition:

$$\mathbb{E}[\eta_{m,f,t}|X_m, k(m), \theta_{m,f,t}] = \boldsymbol{0} \tag{18}$$

which can be thought of as a "control function" or "model based" approach. The second approach imposes the weaker condition that (controlling for type) unobserved determinants of skill outcomes are exogenous with respect to only policy variation:[9]

$$\mathbb{E}[\eta_{m,f,t}|Z_m, k(m), \theta_{m,f,t}] = \boldsymbol{0}. \tag{19}$$

Future sections describing estimation refer to these as the "strict" moment conditions. In either case, identification is guaranteed by rank conditions on the instruments formed from the set of exogenous variables. Appendix D.5 provides details on how these moment conditions are formed in a way that accounts for the fact that neither skills nor type are perfectly observable.[10]

---

[9]Controlling for type in this case is not essential, but allowing for interaction between policies and type strengthens the instruments, and makes the approach robust to any systematic correlation between policies and type.

[10]Since panels available in the PSID are quite long, model predictions of each mother's latent type $k(m)$ are typically quite precise. Using the logic of Bonhomme and Manresa (2015), one could simply assign each mother to their most likely type under the assurance that any classification bias would disappear asymptotically with T. Appendix D.5 clarifies how the moment conditions can be used without this explicit classification using estimated posterior weights over type.

4.3 Estimation of First Stage Parameters

A maximum likelihood approach mobilizes the first stage identification argument for reduced form preferences and distributional parameters. The full log-likelihood is:

$$\mathcal{L}(\mathbf{X}, \Theta) = \sum_{m=1}^{M} \log \left(\sum_{k} \Pi_K(k) \sum_{\varepsilon_m^T} \Pi_W^o(\varepsilon_1) \prod_{t=1}^{T_m} \Pi_W(\varepsilon_{t+1}|\varepsilon_t) \right.$$

$$\times \prod_{d=1}^{D} P_d(S_{m,t}; \Theta)^{\mathbf{1}\{D_{m,t}=d\}}$$

$$\times \left[\frac{1}{\sigma_W} \phi \left(\frac{\log(W_{m,t}^o) - \gamma_{W,0,k} - \gamma_{W,1} Age_{m,t} - \varepsilon_t}{\sigma_W} \right) \right]^{\mathbf{1}\{H_{m,t}>0\}}$$

$$\left. \times \left[\frac{1}{\sigma_\tau} \phi \left(\frac{\log(\tau_{m,t}^o) - \log \left(\frac{\alpha_{\theta,k} \sum_{f \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f,t)}}{\alpha_l + \alpha_{\theta,k} \sum_{f \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f,t)}} (112 - H_{m,t}) \right)}{\sigma_\tau} \right) \right]^{\mathbf{1}\{y(m,t) \in \{1997, 2002\}\}} \right). \quad (20)$$

where $\phi(\cdot)$ is the density of the standard normal. For each observational unit m the likelihood sums over all potential sequences of hidden states and calculates the joint likelihood of each sequence (line 1), the observed choice in each period given observed and unobserved states (line 2), each observed wage (line 3), and each observation of time investment (line 4). The expectation maximization algorithm provides a tractable approach for maximizing the likelihood. Appendix D.4 provides technical details on the implementation of the algorithm. To estimate the model, $\mathcal{E}$ (the number of wage grid points) is set to 5, and Π_W is specified in terms of a single parameter, π_W. For interior grid points, ε_t can move either up or down the grid space with probability $(1 - \pi_W)/2$. For the highest (lowest) grid points, ε_t can only move down (up) with probability $(1 - \pi_W)/2$. A fixed effects regression of log wages on age gives a residual with standard deviation of 0.8. Accordingly, the grid space for ε is set by taking 5 evenly spaced grid points between -1.6 and 1.6 (two standard deviations of the fixed effects residual above and below zero). An analysis of the the Bayesian (BIC) and Akaike Information Criterion (AIC) suggests that $K = 10$ is an appropriate choice for the number of types (see Figure 13).

Table 3 reports the estimates of parameters from this first stage procedure. It documents a substantial degree of heterogeneity across types in all dimensions: the weight on child skills (α_θ), costs of work (α_H), costs of welfare participation (α_A), and productivity in the labor market ($\gamma_{W,0}, \gamma_{W,1}$). Of particular note is the fact that there is no monotonic relationship between any pairs of these latent parameters, highlighting the importance of allowing for unspecified structure in latent heterogeneity. Section 4.5.1 will later explore how these differences in primitive parameters manifest in very different behavior and responses to policy changes, while Section 4.5.3 documents significant heterogeneity in the responsiveness of labor supply at the extensive margin.

The estimation procedure also produces two types ($k = 6$ and $k = 10$) who essentially never participate in welfare programs. Although costs of participation (α_A) must clearly be very large to rationalize this behavior, the likelihood only very weakly distinguishes among these large values.

Table 3: First Stage Estimates

	Global Parameters					
α_C	α_l	α_{WR}	σ_H	π_W	σ_W	σ_τ
1.82	0.95	1.34	0.92	0.64	0.67	0.18
(0.08)	(0.21)	(0.14)	(0.05)	(0.02)	(0.005)	(0.004)

	Type-Specific Parameters					
	α_θ	α_H	α_A	$\gamma_{W,0}$	$\gamma_{W,1}$	Π_K
$k=1$	0.58	2.02	0.37	2.33	0.06	0.03
	(0.19)	(0.22)	(0.22)	(0.87)	(0.03)	(0.01)
$k=2$	0.77	-0.11	5.36	-0.17	0.16	0.04
	(0.25)	(0.22)	(0.29)	(0.49)	(0.02)	(0.01)
$k=3$	0.88	2.63	4.5	3.7	0.07	0.11
	(0.22)	(0.28)	(0.25)	(0.38)	(0.01)	(0.03)
$k=4$	0.79	0.09	2.74	3.42	0.07	0.17
	(0.18)	(0.15)	(0.15)	(0.21)	(0.01)	(0.02)
$k=5$	0.94	1.35	6.55	4.96	0.03	0.24
	(0.21)	(0.2)	(0.31)	(0.21)	(0.01)	(0.03)
$k=6$	1.3	3.3	17.85	4.7	0.05	0.17
	(0.3)	(0.25)	-	(0.31)	(0.01)	(0.03)
$k=7$	3.13	-2.24	2.33	4.45	0.06	0.04
	(0.87)	(0.56)	(0.31)	(0.63)	(0.02)	(0.01)
$k=8$	0.88	-0.31	1.07	1.47	0.17	0.07
	(0.22)	(0.23)	(0.16)	(0.48)	(0.02)	(0.01)
$k=9$	1.54	1.96	2.24	4.04	0.06	0.06
	(0.37)	(0.24)	(0.23)	(0.43)	(0.01)	(0.01)
$k=10$	2.71	-0.24	25.71	4.42	0.06	0.07
	(0.7)	(0.54)	-	(0.49)	(0.02)	(0.02)

This table reports maximum likelihood estimates of the first stage model parameters using the EM-VFI algorithm described in the text. Standard errors are reported in parentheses and are calculated using the typical MLE formula (the inverse of the covariance of the score equations). For two parameters, $\alpha_{A,6}$ and $\alpha_{A,10}$, rank issues make standard errors numerically unstable and these are not reported. See the main text for discussion of this issue.

Since any large value of this parameter will rationalize the data almost equally well, the standard errors are numerically unstable (i.e. very large). For ease of interpretation these standard errors are omitted from Table 3 and the costs are treated as known in future simulation exercises. Fortunately, this makes no difference to observed behavior so long as no policy is considered that somehow increases welfare participation for these types in any substantial way. In practice, the estimated costs are so large that participation is effectively zero in all counterfactuals considered.

4.4   Estimation of Production Parameters

Given available sample sizes it is not feasible to estimate completely age-dependent factor shares. Instead, the factor shares $(\delta_{x,a}, \delta_{\tau,a})$ differ by three stages: (1) Ages 0-5 (early childhood); (2) Ages 6-12 (middle childhood); and (3) Ages 13 and older (adolescence). Appendix D.5 derives a set of instruments based on the control function condition (18) (call these the "model" instruments) to form a moment condition:

$$\mathbb{E}[g_{m,f,t}^{\text{model}}(\boldsymbol{\delta})] = 0.$$

Identically, the weaker restriction (19) provides a set of "strict" instruments with the moment condition:

$$\mathbb{E}[g_{m,f,t}^{\text{strict}}(\boldsymbol{\delta})] = 0.$$

Table 4 reports generalized method of moments estimates (GMM) using both conditions (standard errors are reported in parentheses). Estimates incorporate the scaling coefficients $(\gamma_{\mathcal{Y},C}, \gamma_{\mathcal{Y},B})$ and so each coefficient is the net present value increase (in year 2000 $100,000s USD) in total economic resources of a log unit increase in that input. For example, the point estimate of $\delta_{x,C,0-5}$ when using the control function condition (row 1, column 1) is 0.16, implying that a unit log increase in income results in an increase in cognitive skills for children younger than 5 that is worth $16,000 in future economic resources. This point estimate is different from zero at a 95% significance level. Figure 2 presents the same estimates in graphical format with 95% confidence intervals.

Notice that regardless of whether only policy variables Z_m (the strict IV condition) or all observed outcomes X_m (the control function condition) form the instruments, both estimators approximately agree in terms of point estimates. The estimators also agree on which estimates are significantly different from 0 (at 95% significance) with the sole exception of the factor share of time for cognitive skills $(\delta_{\tau,a,C})$ when children are younger than 5. The similarity of point estimates validates the use of all model instruments for the preferred estimation procedure, which imposes some theoretical content of the model (that factor shares be positive) and rules out implausibly high values of production shares using conservative and weak priors. The latter helps with issues of precision that are known to plague instrumental variables estimators. Define the quasi-likelihood

L_M as

$$L_M(\boldsymbol{\delta}) = -\frac{M}{2} \left(\sum_{m=1}^{M} \sum_{f} \sum_{1997,2002} g_{m,f,t(m,1997)}^{\text{model}}(\boldsymbol{\delta}) \right)' \mathbf{W}^* \left(\sum_{m=1}^{M} \sum_{f} \sum_{1997,2002} g_{m,f,t(m,1997)}^{\text{model}}(\boldsymbol{\delta}) \right)$$

where $\mathbf{W}^*$ is the inverse of a consistent estimate of the covariance of the moment conditions. Chernozhukov and Hong (2003) demonstrate that this function can be used as an approximate log-likelihood because the moments themselves are asymptotically normally distributed. Letting π define priors over the parameters, a Markov Chain Monte Carlo (MCMC) routine can sample from the posterior density

$$p(\boldsymbol{\delta}) = \frac{\exp(L_n(\boldsymbol{\delta}))\pi(\boldsymbol{\delta})}{\int \exp(L_n(\boldsymbol{\delta'}))\pi(\boldsymbol{\delta'})d\boldsymbol{\delta'}}$$

from which the sample mean defines a consistent quasi-Bayesian estimator $\hat{\boldsymbol{\delta}}$ and the quantiles define consistent credibility intervals. Table 4 and Figure 2 also report these preferred estimates with standard errors of the posterior distribution and 95% credibility intervals. While values of factor shares below zero are ruled out, the estimator also significantly moderates the magnitude of the gmm estimates that were found to be positive and significant. Notice also the increase in precision as a result of imposing weak priors.

In order to better understand the role played by priors in this estimation procedure, Figure 14 in Appendix A compares the prior and posterior distributions of each factor share for time and money. In some cases, such as for the share of time in producing behavioral skills for young children, the data do not provide much additional information and the posterior does not shift from the prior. Figure 14 also shows that in most cases the posterior estimate of the factor share is larger in magnitude than what is implied by the prior. This vindicates in some sense the assertion that the priors imposed on this procedure were conservative in nature. Section 4.6.1 will further analyze the economic significance of this preferred set of estimates and benchmark them against related results in the empirical literature.

4.5 Analysis of Behavioral Parameters

4.5.1 Model Fit

To assess fit of the model, Figure 3 compares the mean rates of labor force participation (LFP) and program participation (AFDC) for each latent type[11]. The model exhibits excellent fit of the time series in behavior for each type. As discussed in the introduction, the years between 1990 and 2002 saw quite dramatic changes to the policy landscape and labor market conditions for single mothers, which provides important variation to identify the behavioral parameters of the model.

---

[11]Conditional means by type can be calculated using the posterior weights for each mother-year observation that are produced in the expectation step of the EM algorithm.

Table 4: Cobb-Douglas Share Estimates

| Input | Age | Cog Skills | | | Behav Skills | | |
		IV - Model	IV - Strict	Quasi-Bayes	IV - Model	IV - Strict	Quasi-Bayes
Money	0-5	0.16	0.15	0.05	-0.05	-0.01	0.04
		(0.07)	(0.09)	(0.04)	(0.04)	(0.07)	(0.03)
Money	6-12	-0.01	-0.01	0.01	0.03	0.04	0.02
		(0.02)	(0.03)	(0.01)	(0.02)	(0.02)	(0.02)
Money	13+	0.05	0.06	0.03	-0.01	-0.02	0.01
		(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.01)
Time	0-5	1.22	0.95	0.19	-0.59	-0.72	0.08
		(0.73)	(0.91)	(0.29)	(0.53)	(0.77)	(0.12)
Time	6-12	-0.15	-0.11	0.07	0.46	0.62	0.38
		(0.2)	(0.27)	(0.08)	(0.17)	(0.22)	(0.2)
Time	13+	0.51	0.52	0.19	0.01	0.02	0.24
		(0.19)	(0.23)	(0.15)	(0.19)	(0.22)	(0.12)
Cog Skill	All	0.94	0.92	0.94	0.02	0.02	0.02
		(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Behav Skill	All	0.01	0.0	0.01	0.91	0.91	0.91
		(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)

This table presents point estimates of Cobb-Douglas shares in the technology of skill formation for each skill output, input, and age group using the three strategies outlined in Section 4.4. The column "IV - Model" reports estimates from the GMM procedure using moment conditions described in (50), "IV-Strict" reports estimates from the moment conditions using only policy variation as in (51), while "Quasi-Bayes" reports estimates from the MCMC procedure. Standard errors are reported in parentheses. See Appendix D.5 for more details.

Figure 2: Cobb-Douglas Share Estimates

This figure shows the estimates of the Cobb-Douglas shares of time and money inputs for different ages and skill outputs. "IV - Model" refers to the GMM estimates using moment conditions described in (50), "IV-Strict" refers to estimates from the moment conditions using only policy variation as in (51), while "Quasi-Bayes" reports estimates from the MCMC procedure. Error bars show 95% confidence and credibility intervals. Credibility intervals for the quasi-bayesian estimates are calculated using quantiles from the sampled posterior distribution. See Appendix D.5 for more details.

Figure 3: Model Fit by Latent Type for Program Participation and Labor Force Participation

This figure shows the model's fit of mean rates of labor force participation (LFP) and welfare participation (AFDC) by latent type. Means from the data are calculated using posterior weights generated by the model's estimates. Figure 15 in Appendix A shows the model fit for earnings by type.

A complementary observation is that the panel dimension of the data clearly unveils important degrees of heterogeneity among the population of interest. We see both persistent differences in propensities to work and participate in welfare programs, but also differences in the responsiveness of individuals to these programs. Close inspection of Figure 3 leads to many interesting comparisons across types, and shows that the data support the existence of types with higher rates of both work and welfare use (types 3, 4, 8, and 9), types with high rates of work and low rates of welfare use (types 5, 6, and 10), types with high responsiveness to policy (types 1, 3, 8, and 9) and types with little responsiveness to policy (types 5 and 10). It is important to consider these differences when assessing policy outcomes.

Both propensities and elasticities give important shape to the welfare implications of safety net reforms, and heterogeneity in these behavioral features forms the basis of substantial variation in gains and losses from any reform. To further understand this latent heterogeneity, the next section devotes some time to analyzing elasticities implied by the model.

### 4.5.2 Characteristics by Type

Figure 16 depicts estimated averages of specific demographic characteristics by type. It shows that there are indeed systematic correlations between latent types and fertility behavior (as measured by mother's age at first birth and total births), as well as education (as measured by High School and College graduation). Importantly however, these relationships are not monotonic in the endogenous variables of interest (labor force and program participation) which underscores the importance of controlling beyond observable heterogeneity.

It is also possible to use the model's estimates to examine which latent characteristics determine self-sorting into welfare programs, as in van Dijk (2019). Three parameters in particular are natural candidates as potential drivers of self-sorting into welfare: (1) the cost of work (α_H); (2) the cost of program participation (α_A); and (3) potential earnings. Figure 17 plots these type-specific characteristics against averages of program participation for each type. The estimates suggest that the cost of program participation has the strongest relationship with self-selection, followed by potential earnings. The cost of work does not appear to play a systematic role. The strong selection on unobservable gains that is demonstrated by Figure 17 also emphasises the importance of accounting for unobservable heterogeneity when analyzing work and program participation behavior.

4.5.3 Elasticities

As is typical when designing tax incentives (Saez, 2001), the responsiveness of individual work behavior to incentives will emerge as a crucial determinant of the shape of optimal policies. Discrete choice models typically do not admit constant elasticities, as is often assumed in tax policy analysis (see Saez (2002) for a closely related example). In this model they will vary based on almost every feature of the problem. Table 4 plots average labor supply elasticities for each type for two years,
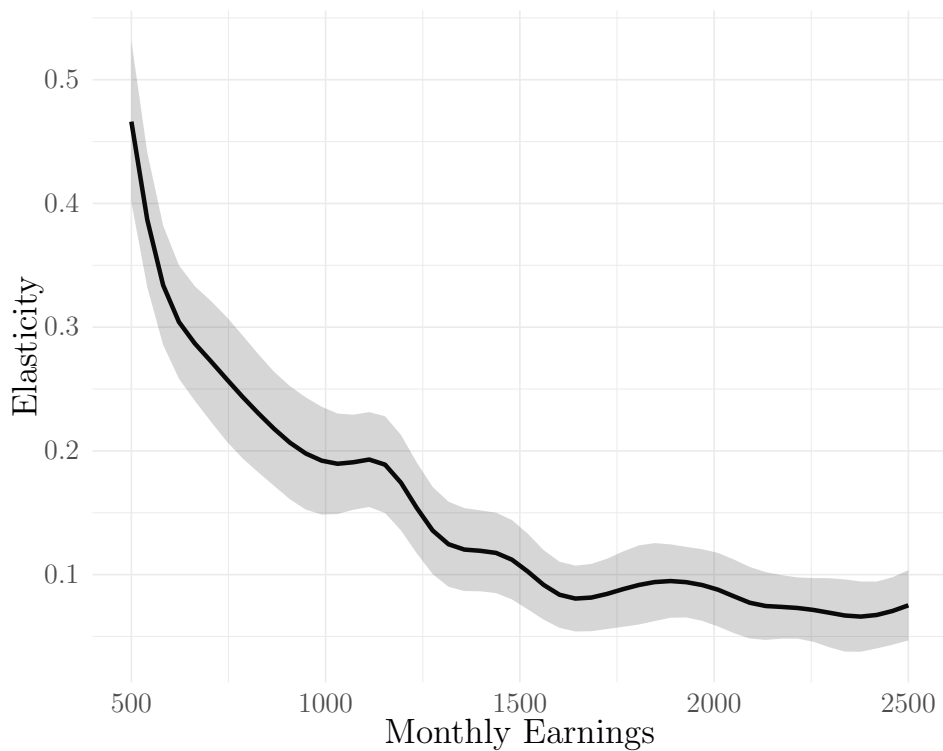
Figure 4: Elasticities by Type

This figure shows the average elasticity of employment for each type, calculated using the sample in 1997 and 2000. Elasticities are calculated for each mother in each potential latent state by increasing post-tax earnings by 10% and dividing the percentage increase in the probability of working by 10%. The average is calculated by taking a weighted mean using the posterior probabilites $q(\cdot|X_m, \hat{\Theta})$ implied by the model likelihood over each latent state.

1996 and 2000. The comparison over years is instructive because the intervening years saw large reductions in welfare participation and large increases in work for some types. Accordingly, Figure 4 delivers a picture with almost as much variation across time periods as across types. The substantial variation in labor supply responsiveness across types echoes other recent results (Attanasio et al., 2018).

For designing optimal nonlinear transfers, the average responsiveness in work behavior at different levels of potential earnings is what matters most. Accordingly, Figure 5 calculates the average labor supply elasticity at each potential earnings for all sample members in the year 2000. It demonstrates that elasticities are highest at the bottom of the distribution of potential earnings, and reduce almost monotonically over the distribution. With the caveat that elasticities in this model are endogenous to policy, bearing this pattern in mind will help the reader to better understand future optimal tax calculations in Section 5.2.

Figure 5: Elasticities by Monthly Wages

This figure shows the average elasticity of employment for each potential earnings level using the sample in 2000. Elasticities are calculated for each mother in each potential latent state by increasing post-tax earnings by 10% and dividing the percentage increase in the probability of working by 10%. The average at any potential earnings level is calculated by taking a weighted mean over each these elasticities using posterior probabilities $q(\cdot|X_m, \hat{\Theta})$ and a Gaussian kernel with a bandwidth of \$100.

4.6 Analysis of Production Parameters

4.6.1 Benchmarking Production Parameters

How do the preferred estimates from Section 4.4 compare to benchmarks in the literature? For cognitive skills, Dahl and Lochner (2012) find that an additional $1000 in annual household income from EITC expansions leads to about a 6% of a standard deviation increase in math and reading scores two years later. For behavioral skills, Akee et al. (2018) find that transfers from casino openings, which increased annual household income by $3500 also lead to a 20% of a standard deviation reduction in behavioral problems.

There is no simple way to compare this paper's estimates to those studies due to the nonlinear effect of income implied by the outcome equation, differences in samples across studies (including the age composition of children and the distribution of income), and the fact that while the model predicts that income and labor supply are simultaneously endogenous, these papers do not control for endogenous responses in labor supply. In order to make some comparison, Figure 18 offers two calculations from the model. The first is a ceterus paribus effect: the average skill impact of an extra $1000 in household income using the distribution of households in the year 2000 as a benchmark, fixing all other choices. The second is an endogenous effect: the average skill impact of increasing household income by $1000 for all choices and allowing decisions to optimally adjust. Appendix D.5.3 provides more details on each calculation. Figure 18 in Appendix A suggests that production parameter estimates in this paper are more conservative than benchmarks from the literature, although there is some overlap in confidence sets.

Bernal and Keane (2011) provide a benchmark for δ_τ by using welfare reform rules as instruments for labor supply to estimate the effect of maternal full-time work in the first five years on cognitive outcomes Bernal and Keane (2011) find that one full year of full-time work leads to an 11% of a standard deviation reduction in cognitive test score outcomes. Since production estimates in this paper are derived from the baseline model, which features a single work decision of 30 hours a week, an equivalent ceterus paribus calculation can be made for a child at age a as $\delta_{\tau,a,C} \log\left(\frac{82}{112}\right)$. Figure 19 in Appendix A reports these estimated ceterus paribus effects in comparison to Bernal and Keane (2011). While point estimates in the model appear to be lower than that found in Bernal and Keane (2011), there is clearly substantial overlap in credibility and confidence intervals for all four parameters.

4.6.2 The Developmental Cost of Work

The relative magnitudes of δ_x and δ_τ will determine whether the decision to work has positive or negative impacts on child development outcomes. Since skils are anchored to the net present value of economic resources, one can calculate a "total" resource effect for a child of age a with a mother

Figure 6: The Total Developmental Effect of Employment

This figure shows the estimated overall effect of single mothers' employment on child human capital as a function of potential earnings, as in Equation 21. The average of this equation is taken over mothers in the sample in the year 2000. 95% credibility intervals are shown, and are calculated by bootstrap sampling from the MCMC chain generated in Section 4.4.

m who makes participation choice p and works 30 hours a week as:

$$(\delta_{x,a,C} + \delta_{x,a,B}) \log \left(\frac{\mathbf{Y}_m(p,1)}{\mathbf{Y}_m(p,0)} \right) + (\delta_{\tau,a,C} + \delta_{\tau,a,B}) \log \left(\frac{72}{112} \right) \tag{21}$$

This calculation offers one way to compare the magnitudes of δ_x and δ_τ, even though the income function $\mathbf{Y}_m$ is clearly not policy invariant. Figure 6 shows the average calculation of the effect of work using equation (21), taking the average for $\mathbf{Y}_m$ over mothers in the sample in the year 2000, assuming both participation ($p = 1$) and non-participation ($p = 0$). The results are quite striking and have important implications for the remaining policy and counterfactual analysis. Figure 6 suggests that the overall effect of work on the value of future child human capital is negative at most earnings levels, with varying levels of uncertainty that depend on the age of the child. While these calculations are made using the preferred set of estimates from the Quasi-Bayesian procedure, Figure 20 in Appendix A repeats the calculations using estimates from the strict instrumental variables condition and shows similar (albeit more imprecise) results. The fact that time appears to "outweigh" money in this outcome equation is a stark prediction that has the potential to significantly dampen the planner's desire to incentivize employment.

4.6.3 The Importance of Accounting For Heterogeneity

The moment conditions described in equation (50) rely on adequately controlling for heterogeneity in the data. The weaker moment conditions that use only policy variation (equation (51)) also rely on the presence of this heterogeneity in that the effects of policies (the instruments) vary by type. Greater heterogeneity in the response to policies therefore provides greater variation for estimation. This section considers the robustness of production estimates to specifications with insufficient heterogeneity. Figure 21 in Appendix A compares estimates for the baseline case ($K = 10$) to estimates from a specification with limited heterogeneity ($K = 2$). While there are no statistically significant differences in pairwise comparisons of estimates, Figure 21 does indicate some substantial differences in point estimates for younger children, and for older children's cognitive skills. Finally, estimates for this alternative specification are uniformly less precise. Overall, the results suggest it is important to adequately control for heterogeneity when estimating production parameters.

4.7 Robustness and Specification Tests

Section 2.9 pointed out that the model admits some fairly strong restrictions on behavior in the data, some of which was directly addressed by an appeal to existing evidence. Section 3 addressed other aspects of the model by looking directly at the data in Table 1. This section provides a more rigorous examination of the model by deriving and conducting three specific empirical tests, with technical details found in Appendix D.6.

First, one of the more provocative aspects of the model is that investments are independent of the child's current stock of skills. This feature allows for drastic simplification of the model's solution properties, but also places strong restrictions on the data. Appendix D.6 shows that the model admits a sibling-pair design: within-household variation in skills should have no predictive power for within-household allocations of time investment. A regression analysis in Appendix D.6 reports coefficients on skills in such a specification that are fairly robustly and reasonably precisely centered around zero, providing some supporting evidence for this feature of the model.

Second, if single mothers were to enjoy wage returns to experience, this would meaningfully change social trade-offs when considering optimal work incentives. It also adds important dynamic empirical content to the model in the form of anticipatory effects from policy changes. Appendix D.6 shows that if this modeling assumption is correct, then labor market experience has no ability to predict wages once other information is accounted for. The results indicate strong support for the assumption of no returns. This is not a general finding but is specific to the individuals in this sample, who are largely characterized by low education. Blundell et al. (2016) similarly find that there are no returns to experience for women with lower levels of education in UK data.

Third, equation (12) outlines a linear dependence structure of child outcomes on income and hours at home, with uniform ceterus paribus effects of work on children. It is reasonable to ask

whether there may be more elaborate substitution patterns in the data. Perhaps higher income individuals can substitute time for money investment, for example. Some evidence suggests that while maternal employment does reduce time with children, it does not reduce time spent in activities that are likely to impact development (Bastian and Lochner, 2021; Hsin and Felfe, 2014). The model rules out such patterns of substitution by predicting that time investment is a constant fraction of time at home. Appendix D.6 tests the model's implications that the fraction of time spent at home is invariant to labor supply and income, finding no statistically significant evidence that this restriction is violated.

5   Designing Optimal Transfers

The model argues that changes in tax and transfer policies can have important direct and indirect impacts on child skill formation through their effects on households' time and money resources. Estimates verify that both time and money play a quantitatively relevant role in skill formation. This section clarifies the immediate implications for the design of transfers by analyzing a simple planner's problem in the spirit of Mirrlees (1971), and showing how the presence of technology parameters (δ_x, δ_τ) changes the calculus of an otherwise classic problem considered first by Diamond (1980). The benchmark case of interest is a universal system of transfers defined by a *retention function* that replaces the existing set of cash assistance policies:

$$\boldsymbol{y}(e) = e - \boldsymbol{\tau}(e)$$

where e is the individual's realized earnings and $\boldsymbol{\tau}$ is the tax function. In the presence of the new policy, the mother's decision problem simplifies. There is no longer a participation decision (cash assistance is universal) and there are no time limits on receipt. The relevant state variables for decision-making reduce down to the triple $(e, \boldsymbol{a}, k)$ where $\boldsymbol{a}$ is the vector of ages of all developing children in the household and

$$e = \gamma_{W,0,k} + \gamma_{W,1,k} Age + \varepsilon$$

is potential earnings. It will be useful to define $s = (k, \boldsymbol{a})$ separately from e. Choices simplify to a work decision indexed by $d \in \{0, 1\}$. Indirect utility is:

$$u_d(y, s) = \tilde{\alpha}_C(s) \log(y) + \tilde{\alpha}_l(s) \log(112 - d \cdot H) - \alpha_{H,k} \cdot d$$

where H is the number of hours worked if working.[12] All remaining dynamics are contained in the recursive coefficients $\tilde{\alpha}_C$ and $\tilde{\alpha}_l$ found in indirect utility (11).[13] The probability of working is:

$$P(s, e) = \left(1 + \exp\left(\frac{u_0(\boldsymbol{y}(0), s) - u_1(\boldsymbol{y}(e), s)}{\sigma_H}\right)\right)^{-1}.$$

---

[12]In the model, $H = 30$.

[13]It may help to remind the reader that $\tilde{\alpha}_C(s) = \alpha_C + \alpha_{\theta,k} \sum_{a \in \boldsymbol{a}} \Gamma_{x,a}$ and $\tilde{\alpha}_l(s) = \alpha_l + \alpha_{\theta,k} \sum_{a \in \boldsymbol{a}} \Gamma_{\tau,a}$.

5.1 Theory

5.1.1 Writing the Planner's Objective

Consider a government planner who chooses household allocations of net income as a function of earned income ($\boldsymbol{y}$) in order to maximize a weighted average of household utilities for this population subject to a constraint on the net present value of economic resources. Let $\boldsymbol{\lambda}$ be the lagrange multiplier on the resource constraint. Since cash assistance policies for this population are only a component of a larger public finance decision, this section poses the planner's problem with a fixed $\boldsymbol{\lambda}$. The multiplier $\boldsymbol{\lambda}$ is often interpreted as the marginal value of public funds and provides a simple way to index (up to scale) the planner's tastes for distribution toward this population (Saez, 2002).

Because allocations of net income and mothers' work decisions both have an effect on the future human capital of children in the household, the planner must account for these effects in their design problem. The model admits a simple form of this accounting, owing to the log-additive contribution of investments to future skills. Consider the effect of a marginal increase in net income for a child at age a on skills at adulthood. Given the linear investment equations described in (9)-(10), this marginal effect (across both skills) has a constant elasticity:

$$\frac{\partial \log(\theta_{18})}{\partial \log(y)} = \delta_\theta^{17-a} \delta_{x,a}.$$

In the next immediate period, the return to a unit log incease in net income is given by the vector of Cobb-Douglas shares for money, $\delta_{x,a}$. This return is compounded into skills in the next period through the matrix of Cobb-Douglas shares for skills, δ_θ, and so on until maturity. Since the net present value of skills is linear in $\log(\theta)$, the marginal return to $\log(y)$ in net present value terms is simply:

$$\gamma_y' \beta^{17-a} \delta_\theta^{17-a-1} \delta_{x,a}.$$

Naturally, a similar expression can be written for the ceterus paribus effect in net present value of a mother's decision to work:

$$\gamma_y' \beta^{17-a} \delta_\theta^{17-a-1} \delta_{\tau,a} \log\left(\frac{112 - H}{112}\right).$$

The planner uses these two expressions to financially account for the effect that policies today have on future skills. For a mother in state s it is useful to define the coefficients:

$$\tilde{\delta}_x(s) = \gamma_y' \sum_{a \in \boldsymbol{a}} \beta^{17-a} \delta_\theta^{17-a-1} \delta_{x,a} \tag{22}$$

$$\tilde{\delta}_\tau(s) = \gamma_y' \sum_{a \in \boldsymbol{a}} \beta^{17-a} \delta_\theta^{17-a-1} \delta_{\tau,a}. \tag{23}$$

Using this simplified accounting for skill formation one can conveniently write the planner's problem

as:

$$\max_{\boldsymbol{y}} \sum_{s,e} \pi(s,e) \left[\mu(s,e) \max_{d \in \{0,1\}} \{u_d(\boldsymbol{y}(d \cdot e), s) + \epsilon_d\} \right.$$

$$+ \boldsymbol{\lambda}(1 - P(s,e)) \left[\tilde{\delta}_x(s) \log(\boldsymbol{y}(0)) + \tilde{\delta}_\tau(s) \log(112) - \boldsymbol{y}(0) \right]$$

$$\left. + \boldsymbol{\lambda} P(s,e) \left[\tilde{\delta}_x(s) \log(\boldsymbol{y}(e)) + \tilde{\delta}_\tau(s) \log(112 - H) + e - \boldsymbol{y}(e) \right] \right] \quad (24)$$

where $\pi(s,e)$ is the distribution of individuals over (s,e) and $\mu(s,e)$ is the Planner's weight on households of type (s,e).[14]

5.1.2 First Order Conditions

Introducing the technology parameters δ_x and δ_τ leads to conceptually substantive changes to the planner's calculus. The first order condition with respect to $\boldsymbol{y}(e)$ is:

$$\sum_s \pi(s,e) \left\{ P(s,e) \left[\frac{\mu(s,e)\tilde{\alpha}_C(s) + \boldsymbol{\lambda}\tilde{\delta}_x(s)}{\boldsymbol{y}(e)} - \boldsymbol{\lambda} \right] \right.$$

$$\left. + \boldsymbol{\lambda} \frac{\partial P(s,e)}{\partial[u_1(\boldsymbol{y}(e),s) - u_0(\boldsymbol{y}(0),s)]} \frac{\tilde{\alpha}_C(s)}{\boldsymbol{y}(e)} \left[e + \boldsymbol{y}(0) - \boldsymbol{y}(e) + \mathcal{D}(s,e) \right] \right\} = 0 \quad (25)$$

where $\mathcal{D}(s,e)$ is the total effect on skills, as measured by future economic resources, of a mother's decision to work in state (s,e):

$$\mathcal{D}(s,e) = \tilde{\delta}_x(s) \log\left(\frac{\boldsymbol{y}(e)}{\boldsymbol{y}(0)}\right) - \tilde{\delta}_\tau(s) \log\left(\frac{112}{112 - H}\right).$$

To build intuition, consider each line in the first order condition. The first line represents the direct marginal benefits and costs of a small increase in net income at earnings e, consisting of the mother's marginal utility of consumption (weighted by $\mu(s,e)$), the marginal return to future resources (scaled by $\tilde{\delta}_x(s)$), and the resource cost of this marginal increase ($\boldsymbol{\lambda}$). Each of these terms is weighted by the fraction of women in each state who work. These are the direct fiscal effects of the marginal change. The second line is the behavioral effect that the planner must consider when marginally increasing $\boldsymbol{y}(e)$. The response in work probabilities is the product of the responsiveness of choices to changes in utility and the marginal change to utility induced by a marginal increase in $\boldsymbol{y}(e)$. The resource cost of this behavioral change is given by the square bracketed term and incorporates costs today as well as the skill effects of changes in work behavior ($\mathcal{D}(s,e)$). This term may be positive or negative depending on the relative importance of time (δ_x) and money (δ_τ) in the formation of skills. An envelope condition guarantees that marginal changes in utility in response to this behavioral change are zero (McFadden, 1981). In terms of traditional public

---

[14]Note that if $\pi(s,e)$ is in stationary distribution, the problem is equivalent to maximizing a weighted average of household's dynamic values.

finance accounting, the introduction of child skill formation makes a novel addition to the direct effect of marginal changes through $\tilde{\delta}_x(s)$ and to the behavioral effect through $\mathcal{D}(s, e)$.

Following identical logic, the first order condition with respect to $\boldsymbol{y}(0)$ is

$$\sum_{s,e} \pi(s,e) \left\{ (1 - P(s,e)) \left[\frac{\mu(s,e)\tilde{\alpha}_C(s) + \boldsymbol{\lambda}\tilde{\delta}_x(s)}{\boldsymbol{y}(0)} - \boldsymbol{\lambda} \right] \right.$$
$$\left. - \boldsymbol{\lambda} \frac{\partial P(s,e)}{\partial[u_1(\boldsymbol{y}(e), s) - u_0(\boldsymbol{y}(0), s)]} \frac{\tilde{\alpha}_C(s)}{\boldsymbol{y}(0)} \left[e + \boldsymbol{y}(0) - \boldsymbol{y}(e) + \mathcal{D}(s,e) \right] \right\} = 0 \quad (26)$$

with a key difference being that the marginal increase in net income when not working induces a small reduction in work behavior.

5.1.3   Optimal Size

Rearranging and combining equations (25) and (26) gives the first key condition, which characterizes the optimal size of cash assistance:

$$\mathbb{E}\left[(1 - P(s,e))\boldsymbol{y}(0) + P(s,e)\boldsymbol{y}(e) \right] = \mathbb{E}\left[\frac{\mu(s,e)\tilde{\alpha}_C(s)}{\boldsymbol{\lambda}} + \tilde{\delta}_x(s) \right]. \quad (27)$$

It states that the overall generosity of the retention function $\boldsymbol{y}$ – as measured by average net income – is equal to the average planner weight on consumption relative to the marginal value of resources (given by $\mu(s,e)\tilde{\alpha}_C(s)/\boldsymbol{\lambda}$) plus the average importance of money for child development ($\tilde{\delta}_x(s)$) across households. It is convenient to define this *effective* weight on households (s, e):

$$\boldsymbol{w}(s,e) = \frac{\mu(s,e)\tilde{\alpha}_C(s)}{\boldsymbol{\lambda}} + \tilde{\delta}_x(s).$$

Relative to a planner with purely redistributive preferences, accounting for the spillover effect on child human capital increases the effective weight and implies that transfers to these households should be larger if money matters for child development. This condition will provide a simple way to quantitatively benchmark the generosity of the optimal program against the generosity of existing programs.

5.1.4   Optimal Shape

Define $\eta(s, e)$ as the semi-elasticity of $P(s, e)$ with respect to the difference in utilities:

$$\eta(s,e) = \tilde{\alpha}_C(s) P(s,e)^{-1} \frac{\partial P(s,e)}{\partial[u_1(\boldsymbol{y}(e), s) - u_0(\boldsymbol{y}(0), s)]}.$$

Additionally, define $\boldsymbol{y}^*$ as the first-best allocation:

$$\boldsymbol{y}^*(e) = \mathbb{E}[\boldsymbol{w}(s,e)|e, d = 1], \qquad \boldsymbol{y}^*(0) = \mathbb{E}[\boldsymbol{w}(s,e)|d = 0]$$

which is the solution to the planner's problem without incentive compatibility constraints (when $\eta(s,e) = 0$).[15] The second key equation rearranges and relabels terms in (25) to characterize the

---

[15]See Appendix E for a derivation of the first-best allocations.

optimal shape of cash assistance in terms of the weights $\boldsymbol{w}$, semi-elasticities, and the net effect of employment on skills:

$$\boldsymbol{y}(e) = \boldsymbol{y}^*(e) + \frac{\mathbb{E}\left[\eta(s,e)\left(\boldsymbol{y}(0) - \boldsymbol{y}^*(e) + e + \mathcal{D}(s,e)\right)\big|e, d = 1\right]}{1 + \mathbb{E}[\eta(s,e)|e, d = 1]}. \tag{28}$$

It states that allocations can be interpreted as a first best plus a wedge term that increases with the responsiveness of individuals to work incentives, $\eta(s,e)$, with potential earnings, e, and with the net effect of work on child skills, $\mathcal{D}(s,e)$. The first two components of the wedge reflect elements of the planner's trade-offs that have been extensively studied. Accounting for child development introduces the third component, and can have a meaningful impact on optimal work incentives. If money dominates time in skill formation, $\mathcal{D}$ takes positive values and the size of the wedge relative to first-best allocations increases as there are additional benefits to incentivizing employment. When time dominates money, the inverse logic holds true.

The formula is complicated by the fact that the conditional distribution of household states s and semi-elasticities $\eta(s,e)$ may change with earnings. To further develop intuition, it helps to consider the case in which the state of all households in the population is set to some s, with weights $\boldsymbol{w}$ and semi-elasticities η held constant. In this case the formula simplifies to:

$$\boldsymbol{y}(e) = \boldsymbol{w} + \frac{\eta}{1+\eta}\left[\boldsymbol{y}(0) - \boldsymbol{w} + e + \mathcal{D}(s,e)\right]$$

with an accompanying expression for average work incentives:

$$\mathbb{E}[\boldsymbol{y}(e) - \boldsymbol{y}(0)|d = 1] = \frac{\eta}{1 + \eta - \mathbb{E}[d]}\mathbb{E}\left[e + \mathcal{D}(s,e)|d = 1\right] \tag{29}$$

that helps to clarify the problem's mechanics. When δ_τ is sufficiently large $\mathcal{D}$ takes negative values and so the planner optimally discourages work by making it less rewarding compared to not working. Furthermore, applying the implicit function theorem gives:

$$\frac{d\boldsymbol{\tau}(e)}{de} = \frac{1 - \tilde{\delta}_x(s)/\boldsymbol{y}(e)}{1 + \eta - \tilde{\delta}_x(s)/\boldsymbol{y}(e)}. \tag{30}$$

In this simplifying case, equation (29) tells us that the relative magnitude of δ_x and δ_τ determines average work incentives, while (30) states that only δ_x matters for how these incentives change with e. If $\delta_x = 0$, the optimal marginal tax rate is $1/(1 + \eta)$, while if δ_x is sufficiently large, the possibility of negative marginal tax rates emerges. In the upcoming quantitative analysis these simplifying assumptions to do not apply, but the insights regarding the effect that skill formation has on optimal policy are borne out regardless.

Optimal Employment Subsidies Prior work (Diamond, 1980; Saez, 2002) has demonstrated that the introduction of an extensive marginal work decision can introduce employment subsidies that are optimal at the bottom of the income distribution. Appendix E.0.1 verifies this property in the standard case without skill formation and provides some additional insight into the implications

of child skill formation for the presence and size of these earnings subsidies. For example, if work subsidies are optimal, then increases in the factor share of money (δ_x) lead to increases in the size of the subsidies. Most importantly, accounting for skill development implies that there are no employment subsidies at the bottom if δ_τ is sufficiently large relative to δ_x. In this case, it may be optimal for the planner to implement work penalties ($\boldsymbol{y}(e) < \boldsymbol{y}(0)$) for individuals whose potential earnings are not sufficiently high. This observation is particularly important for understanding the upcoming quantitative exercises.

5.2 Quantitative Analysis

What is the optimal policy implied by estimates of preferences and the technology of skill formation? What effect does accounting for children's skill formation have on the shape and size of optimal transfers? This section of the paper combines the estimated model with results from the theoretical analysis of Section 5.1 to answer these questions. Policies considered in this section are "optimal" in the sense that they are chosen to maximize the welfare criterion introduced in (24) with weights μ set to be equal across households, assuming a benchmark distribution of mothers over states. This section proceeds in three steps. The next section uses transfers to equivalent households without children to calibrate a measure of $\mu/\boldsymbol{\lambda}$. Then, the optimal generosity of the program is calculated using equation (27) and compared to the estimated generosity of prevailing programs in the benchmark. The final section solves the planner's problem as stated above and conducts comparative static exercises to explore the quantitative influence of children's human capital formation on optimal policy.

Benchmark Economy   For this policy problem, the sample distribution of mothers over states and policies in the year 2000 forms the benchmark economy. The likelihood evaluated at model estimates produces a posterior distribution for each mother m over unobserved states in the year 2000.

5.2.1 Calibrating the Marginal Value of Resources

Solving the maximization problem as phrased above requires a choice of $\mu/\boldsymbol{\lambda}$, the planner's effective weight on single mothers in their resource allocation problem (24). It measures how the planner values a marginal increase in the net income of households in the population of interest (μ) against reallocating those resources somewhere else in the economy ($\boldsymbol{\lambda}$). In this sense, the ratio $\mu/\boldsymbol{\lambda}$ measures the planner's tastes for redistribution toward this population. Equation (27) suggests one approach for choosing the ratio. Assuming that tastes for redistribution are not affected by the presence of children, Equation (27) simplifies to:

$$\mathbb{E}[y(e)] = \frac{\mu \alpha_C}{\boldsymbol{\lambda}}$$

Figure 7: Optimal Generosity vs Actual Generosity of Cash Transfers

This figure shows the optimal generosity of transfers in the model as calculated in (32) compared to their actual generosity in the sample, calculated in (31). These calculations are made using averages across all households (the case of "Any" number of children), and then for households with one, two, and three children. Error bars show 95% credibility intervals, calculated by drawing 100 bootstrap samples from the estimated asymptotic distribution of preference parameters and from the estimated posterior over production parameters using the MCMC chain.

for the same population of women with no children in the household. Thus, the average generosity of non-child related programs (as measured by average net income) provides a measure for the planner's tastes for redistribution. In the model, this is average net income for households that can only participate in food stamps. The implied value for $\mu/\boldsymbol{\lambda}$ is therefore

$$\frac{\widehat{\mu}}{\lambda} = \hat{\alpha}_C^{-1} \mathbb{E}[y(e)|\hat{\Theta}, \text{No Children}].$$

5.2.2 Optimal Generosity

Taking tastes for redistribution as given, Equation (27) offers a way to benchmark the actual generosity of cash transfers with optimal generosity as implied by this formula. Given parameter estimates, the two calculations are:

$$\widehat{\text{Actual Generosity}} = \mathbb{E}[y(e)|\hat{\Theta}] \tag{31}$$

$$\widehat{\text{Optimal Generosity}} = \mathbb{E}\left[\widehat{\frac{\mu\tilde{\alpha}_C}{\boldsymbol{\lambda}} + \tilde{\delta}_x(s)}\right]. \tag{32}$$

The calculations above provide a way to compare the optimal amount of resources allocated to single mothers relative to the actual amount, under the assumption that total resources allocated to the same agents in a counterfactual without children is an appropriate measure of the government's

weight on this population. Using such a measure of effective weights is likely conservative due to the fact that the planner's problem does not consider the utility of children. Factoring this would provide a mechanism through which weights on households with children mechanically increase. Furthermore, it is plausible that innate tastes for distribution to households with children may be larger than those without, which this approach does not factor (Saez and Stantcheva, 2016). Either such mechanism would imply that this paper's approach – which uses implied transfers to households without children to measure weights – under-estimates optimal generosity.

Equation (27) also implies that if transfers can be conditioned on family size, then net resources should be increasing in the number of children in the household. Empirically this is due in main part to the fact that the contribution of household income to future resources, $\tilde{\delta}(s)$, is increasing in the number of children inside the household. This is true regardless of whether investments inside the household are completely private or public.[16] This section therefore explores the quantitative implications of Equation (27) by comparing actual generosity to optimal generosity for all households with children, then separately by households with one, two, and three children. All comparisons are made using the sample distribution of households and state policies in the year 2000, after the most significant changes to welfare and tax policy have occurred. The quantitative results are robust, however, to the choice of year for the exercise.

Figure 7 shows these calculations using all households, and households with one, two, and three children respectively. The headline result is that, overall, households with children receive about $200 less in net resources per month ($2400 per year) than what is considered optimal in the model. When assessing this calculation, it bears repeating that: (1) estimates of δ_x from the model are smaller than a number of headline results from the literature; and (2) the anchoring of skills to future economic resources – which maps δ_x to $\tilde{\delta}_x$ – also took a conservative approach. Overall then, $200 per month should be viewed as a conservative estimate of the difference between actual and optimal generosity. Figure 7 also shows that, while optimal net resources should be increasing in family size, actual net resources are not. The latter result is due to changes in the composition of types by household size, with mothers that have more children tending to face lower earnings and higher costs of participation in the labor market. As a result, the largest deviation between actual resources and optimal resources for households is for the case with three children. Here the difference is about $600 per month.

While Equation (27) provides a clear way to assess the size of transfer programs compared to some optimal benchmark, it offers no guide to how transfers should be allocated within a population. How should incentives be designed, and how does the technology of skill formation affect the solution? The next section solves the Planner's problem under a number of scenarios in order to answer these questions.

---

[16]Appendix C demonstrates this point.

Figure 8: Optimal Policy vs US Average in Year 2000

This graph shows the solution for the optimal non-linear assignent described in Equation (24) using the empirical distribution of households in the year 2000. 95% bootstrapped credibility intervals are shown for this solution, where parameters are drawn from the Markov Chain described in Section 4.4. "US Average w/ TANF" shows average net income for the sample when individuals are participating in cash welfare, while "US Average w/out TANF" shows the average without participation. Dollar amounts are reported in year 2000 USD. The solution is calculated using Newton's Method.

5.2.3 Optimal Nonlinear Transfers

The previous section demonstrated that when planner weights are calibrated to match average transfers to women without children, an optimal system of transfers that properly incorporates the human capital spillovers for children is more generous than what is observed in the data. This section solves the Planner's problem taking as given the distribution of household demographics and types for mothers in the sample in the year 2000 and using these same calibrated weights. The solution is derived over a flexible space of nonlinear policies that allows for a discontinuity at zero.[17]

Figure 8 depicts the solution, showing the optimal assignment of net income as a function of

---

[17]y is approximated as a transfer at $e = 0$ with a separate continuous function over $(0, \infty)$ defined as a quadratic spline with evenly spaced knots at every $100 of monthly earnings up to $1200, with a linear form thereafter. This results in a function defined by 14 parameters which are chosen to maximize (24) using Newton's method. Additional knot points made no further difference to the planner's objective.

earnings with 95% credibility intervals at each point in the schedule. For reference, Figure 8 also plots net average net income as a function of earnings faced by sample members in the year 2000. Since participation is voluntary in this baseline, lines are drawn assuming both participation and non-participation in TANF.

There are two striking features of the optimal policy. First, the optimal assignment features a much more generous transfer for those not earning (i.e. at zero) relative to the empirical benchmark in the year 2000. The difference is about $800 per month relative to households participating in TANF. The effective difference is larger still since the optimal policy is universal, while many in the benchmark economy do not participate in TANF due to time limits and participation costs.

The second striking feature is the drastic reduction in net income for those earning small positive amounts. Both features combine to create quite drastic discentives for labor supply among single mothers with low potential earnings: a result that challenges the stated goals of PRWORA which prioritized changes to incentives that would encourage labor force participation and reduce dependence on cash welfare. The optimal system of transfers seeks to do the opposite for mothers with low potential earnings by making the decision to not work more financially rewarding. The presence of this feature is due entirely to the influence of mothers' time at home for future skill formation. Since time is estimated to be relevant in the production of skills, it is developmentally costly for single mothers to work if they earn very little in the labor market. This was demonstrated in Figure 6. The optimal policy internalizes this by creating a strong discentive for work at these levels of earnings. The remainder of the schedule appears quite similar to the average US benchmark. Of course, with much higher generosity at zero earnings, overall generosity of the optimal program is larger, as the previous section demonstrated.

Section 5.1 emphasized that properly accounting for spillovers in children's human capital can have an important effect on the shape of optimal transfers. The real magnitude of these differences depends on the estimates themselves, which two adjustments to the optimal policy problem can now help to illustrate. They disentangle the separate effects that each component of the technology of skill formation – δ_x and δ_τ – has on the shape of optimal transfers by first re-solving the policy problem with δ_τ set to zero (i.e. ignoring time effects) and then re-solving again with child human capital ignored completely (both δ_τ and δ_x set to zero).

Figure 9 shows the results for the first exercise, with the factor share on time (δ_τ) set to zero. It also depicts for reference the optimal policy from the benchmark case in Figure 8. The contrast between the new solution and the baseline policy demonstrates one of the key takeaways from Section 5.1: the extent of optimal work incentives depends on the relative magnitudes of δ_x and δ_τ. In particular, the optimal policy features work credits (a nominal transfer for earning any positive amount) in cases where δ_τ is sufficiently small. Accordingly, when δ_τ is ignored entirely, the optimal policy features a work credit of approximately $200 per month. More generally, the new policy does not exhibit the drastic work disincentives of the benchmark case, and allocations

Figure 9: Optimal Policy Ignoring Time in Child Human Capital Formation

This graph shows the solution for the optimal non-linear assignent described in Equation (24) when the Cobb-Douglas shares on time (δ_τ) is set to zero. 95% boostrapped credibility intervals are shown for this solution, where parameters are drawn from the Markov Chain described in Section 4.4. The optimal solution from the benchmark case in Figure 8 is also shown. Dollar amounts are reported in year 2000 USD.

Figure 10: Comparison of Transfers Across Scenarios
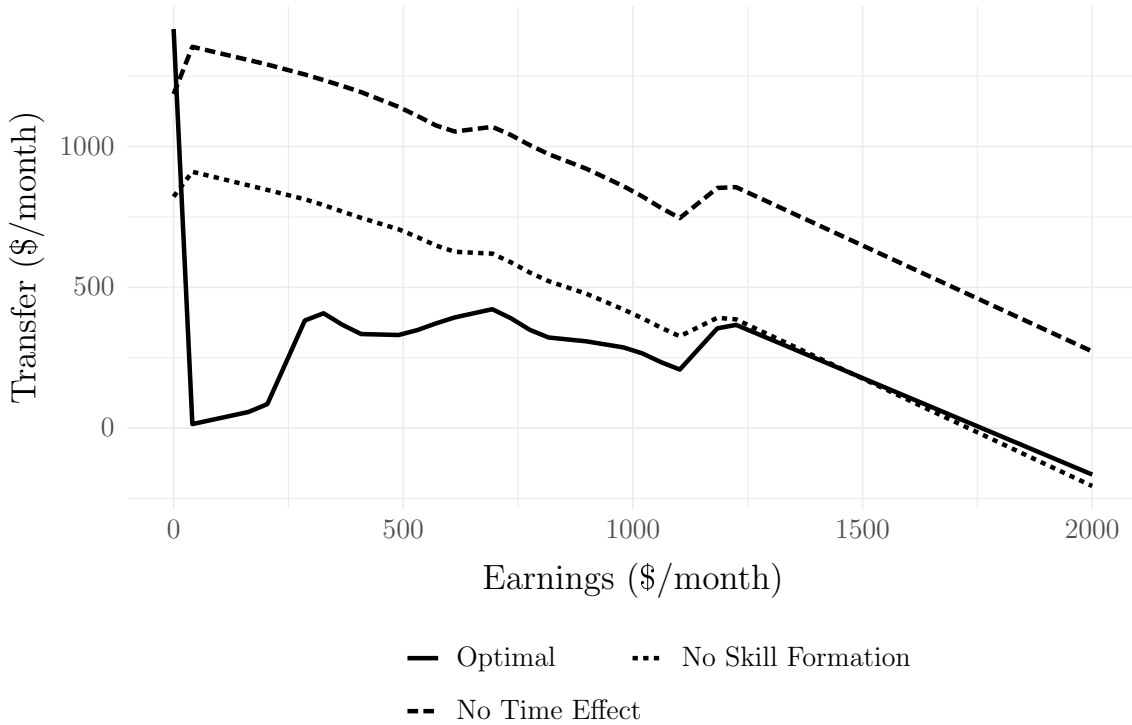
This figure shows the monthly transfer (defined as net income allocation minus earnings) for each of the three optimal policy scenarios: (1) The benchmark case using all parameter estimates; (2) The case in which δ_τ is set to zero; and (3) The case in which both δ_x and δ_τ are set to zero.

are more generous everywhere except for those not earning. While equation (27) requires that average allocations are invariant to δ_τ, this exercise demonstrates the important quantitative role played by δ_τ in determining the shape of optimal policy.

To more clearly compare the shape of optimal transfers, Figure 10 shows the value of monthly transfers as a function of monthly earnings for each of the previous exercises alongside the final exercise in which optimal transfers are derived when child human capital spillovers are ignored completely. Relative to the case in which only δ_x is accounted for in the Planner's objective, the new solution is almost a parallel shift down across the schedule, with a small decrease in the value of the work credit (as suggested by the partial effects studied in Section 5.1.4). Figure 10 suggests that while incorporating the effect of mothers' work on the development of children has a major impact on the shape of optimal transfers, the main impact of incorporating the effect of household income is on the size of optimal transfers.

6   Welfare Reform

The period from 1990 to 2000 saw quite drastic changes to cash transfers for families with children. Most significantly, PRWORA introduced federally mandated time limits on receipt of cash welfare, institutionalized work requirements for welfare participants, and allowed greater leave for state governments to reallocate federal funds across program initiatives. At the same time, several expansions to the Earned Income Tax Credit transformed federal taxes into a major instrument for poverty alleviation. A large literature has sought to estimate and decompose the causal effect of these reforms on welfare participation and labor supply (Chan, 2013; Grogger, 2003; Hoynes, 1996; Meyer, 2002; Keane and Wolpin, 2010). However the overall effects of these important changes to income support on child development outcomes is unkown. This section analyzes a particular policy counterfactual – in which the policy environment is "frozen" in 1996 – to understand the impact of policy changes after this time on maternal welfare and child development outcomes.

By pausing changes to the policy environment from 1996 onward, the counterfactual undoes changes to welfare implementated by PRWORA as well as several of the later expansions to the EITC. In doing so, the model creates a counterfactual calculation of welfare and child outcomes that is not readily estimable from the data. To give a sense of this empirical challenge, Figure 12 in Appendix A reports the estimated effect of exposure to welfare reform on cognitive and behavioral outcomes using within-family variation in exposure. The lack of any conclusive evidence from these regressions makes it clear that without further sources of variation it is difficult to derive empirical conclusions on the effect of welfare reform directly from these data.[18]

Returning to the counterfactual, Table 5 presents the impact on skill outcomes, as well as each type's consumption equivalent valuation (CEV) of the counterfactual: the percentage increase in consumption in every period that would achieve the same average change in values relative to the baseline as the counterfactual policy. Thus, for example, type (3) mothers receive an average increase in welfare from "undoing" welfare reform and later EITC expansions that is equivalent to a 2.42% increase in consumption in every period until their last child matures. Across types, the average welfare gain from the policy counterfactual is equivalent to a 3.35% increase in consumption, however this aggregate result masks substantial heterogeneity. Types (1), (8), and (9) represent about 16% of the population and all enjoy quite substantial welfare gains. On the other hand, types (2), (5), (6), (7), and (10) represent about 56% of the population and all endure smaller but statistically significant welfare losses. In this sense, one of the important impacts of welfare reform can be viewed as a redistribution away from a small number of types that benefited substantially from the old system of transfers to a broader population. These results add a new dimension to the cross-sectional analysis performed by Moffitt (2015), who observes that the overall effect of changes

---

[18]It is worth noting that in this design the effect of welfare reform exposure may be confounded by birth order or cohort effects. Children in the household with higher levels of exposure to welfare reform are by definition older.

to child related transfers has been to target the "deserving poor": low income individuals who nonetheless work. The analysis here confirms that this change results in meaningful redistribution of welfare outcomes across distinct groups of individuals.

Changes in skills are reported separately for cognitive and behavioral skills, with results anchored to their net present value impact on economic resources in thousands of dollars. A third column reports the total effect on skills in net present value. While reversing the welfare reform era does not have a universally statistically significant effect on skill outcomes, the outcomes for all types are positive, and the estimated total net present value effect on skills (per child) amounts to $1,859 in year 2000 USD. As was the case for maternal welfare, these aggregate effects conceal considerable heterogeneity across groups of children, with large skill gains applying to a small fraction of the population, moderate skill gains for some, and statistically insignificant skill losses for others.

In order to better understand the primitives that determine who benefits and who loses from reversing welfare reform, Figure 11 plots the consumption equivalent valuations for each mother against three type-specific primitives: the cost of program participation (α_A), the cost of labor force participation (α_H), and the mean of log wages for each type at Age 25 ("Log Wage"). Figure 11 demonstrates that costs of program participation and labor market productivity appear to be the strongest predictors of welfare gains. The major benefactors of the pre-reform policy environment were single mothers with a higher propensity to participate in AFDC, and who faced lower returns to labor market participation. Figure 23 in Appendix A shows that the exact same pattern holds when plotting total skill impacts against type characteristics, and so the same conclusions apply for the determinants of skill outcomes across types.

What are the consequences for the model's predictions if latent heterogeneity is not adequately accounted for? Table 6 reports results using model estimates of the model when the number of types K is set to $K = 2$. Section 4.6.3 used these estimates to show that properly accounting for latent heterogeneity is important for getting estimates of production parameters right. Accordingly, these calculations use the preferred production parameter estimates from the baseline specification to forecast the impacts on skill outcomes. This simpler specification of heterogeneity is able to capture some of the basic contrasts in how mothers value different policies: type (1) mothers subsantially benefit from undoing welfare reform while type (2) mothers are worse off. However the model predicts much more substantial overall welfare gains from the counterfactual than in the baseline model. The reason for this is intuitive: in the absence of sufficient persistent ex-ante heterogeneity in work and participation behavior, more of this variation is attributed to ex-post risk. Accordingly, this increases the average valuation of policy environments with more insurance against these shocks, which the pre-welfare reform environment provides relative to post-reform. It is well known that the degree of ex-ante heterogeneity is a key object in determining the decomposition of welfare effects into redistribution and insurance components (Bhandari et al.,

Table 5: Impacts of "No Welfare Reform" Counterfactual

	Cog	Behav	Total Skill	% CEV
$k = 1$	7.3	8.41	15.71	74.36
	(3.2)	(2.72)	(4.22)	(8.91)
$k = 2$	0.54	0.76	1.3	-0.76
	(0.26)	(0.34)	(0.51)	(0.43)
$k = 3$	1.02	1.2	2.23	2.42
	(0.48)	(0.39)	(0.65)	(1.04)
$k = 4$	1.39	1.85	3.25	3.12
	(0.74)	(0.61)	(1.02)	(0.69)
$k = 5$	-0.09	0.13	0.04	-2.19
	(0.12)	(0.09)	(0.16)	(0.16)
$k = 6$	-0.27	0.11	-0.17	-2.78
	(0.2)	(0.13)	(0.25)	(0.15)
$k = 7$	-0.01	0.18	0.17	-1.98
	(0.14)	(0.15)	(0.24)	(0.41)
$k = 8$	1.42	1.55	2.97	11.73
	(0.66)	(0.51)	(0.83)	(5.28)
$k = 9$	3.1	3.91	7.01	18.5
	(1.33)	(1.2)	(1.72)	(4.33)
$k = 10$	-0.23	-0.05	-0.28	-3.52
	(0.11)	(0.06)	(0.14)	(0.28)
Total	(0.76)	(1.09)	(1.86)	3.35
	(0.377)	(0.34)	(0.486)	(0.66)

This tables reports the impacts on child skills and maternal welfare of a policy counterfactual in which the policy environment is paused in 1996 prior to the implementation of PRWORA and later expansions in the EITC. Skill impacts are reported in terms of their net present value impact on economic resources in thousands of year 2000 USD. Consumption Equivalent Valuations of the policy are reported as the percentage increase in consumption in every period that would achieve the same average increase in values for each type. Standard errors are reported in parentheses and calculated by bootstrap sampling from the estimated parameter distribution described in Section 4.

Figure 11: Heterogeneous Effects of "No Welfare Reform" on Welfare

This figure plots each type's consumption equivalent valuation of the "No Welfare Reform" policy counter-factual (in which the policy environment is paused in 1996) against each type's cost of program participation (α_A), cost of work (α_H), and mean log wages at age 25 ("Log Wage").

Table 6: Impacts of "No Welfare Reform" Counterfactual With Limited Heterogeneity

	Cog	Behav	Total Skill	% CEV
$k = 1$	-0.51	1.24	0.73	12.54
	(3.44)	(3.3)	(4.98)	(3.33)
$k = 2$	-3.04	-5.36	-8.4	-1.96
	(1.09)	(1.23)	(1.74)	(1.59)
Total	-1.39	-1.07	(-2.46)	7.47
	(2.584)	(2.553)	(3.8)	(2.66)

This tables reports the impacts on child skills and maternal welfare of a policy counterfactual in which the policy environment is paused in 1996 prior to the implementation of PRWORA and later expansions in the EITC. Skill impacts are reported in terms of their net present value impact on economic resources in thousands of year 2000 USD. Consumption Equivalent Valuations of the policy are reported as the percentage increase in consumption in every period that would achieve the same average increase in values for each type. Standard errors are reported in parentheses and calculated by bootstrap sampling from the estimated parameter distribution described in Section 4.

2021; Benabou, 2002; Floden, 2001), which this exercise duly demonstrates. Under-estimating the degree of ex-ante heterogeneity in this environment overstates the insurance value of policy reforms.

In a similar fashion, since more risk is dispersed within types rather than across types, there is no statistically significant effect on child outcomes from the counterfactual. This also contrasts with the baseline model which predicted significant and positive developmental effects for types with persistently lower wages and higher propensities to participate in welfare programs.

Finally, the counterfactual experiment provides an estimate of the effect of welfare reform on work and program participation. Unsurprisingly, Figure 22 in Appendix A confirms that changes in the policy environment from 1996 onwards (including welfare reform and later EITC expansions) led to an increase in labor force participation and a decrease in welfare enrollment, especially among types with high initial rates of AFDC participation. Chan (2013) provides a thorough decomposition of the effects of individual policies on these outcomes using a similar dynamic model without children.

7   Conclusion

The results of this paper demonstrate the substantial effect that properly accounting for future child skill development can have on the design of an otherwise standard policy aimed at providing

Figure 12: Effects of Exposure to Welfare on Child Outcomes
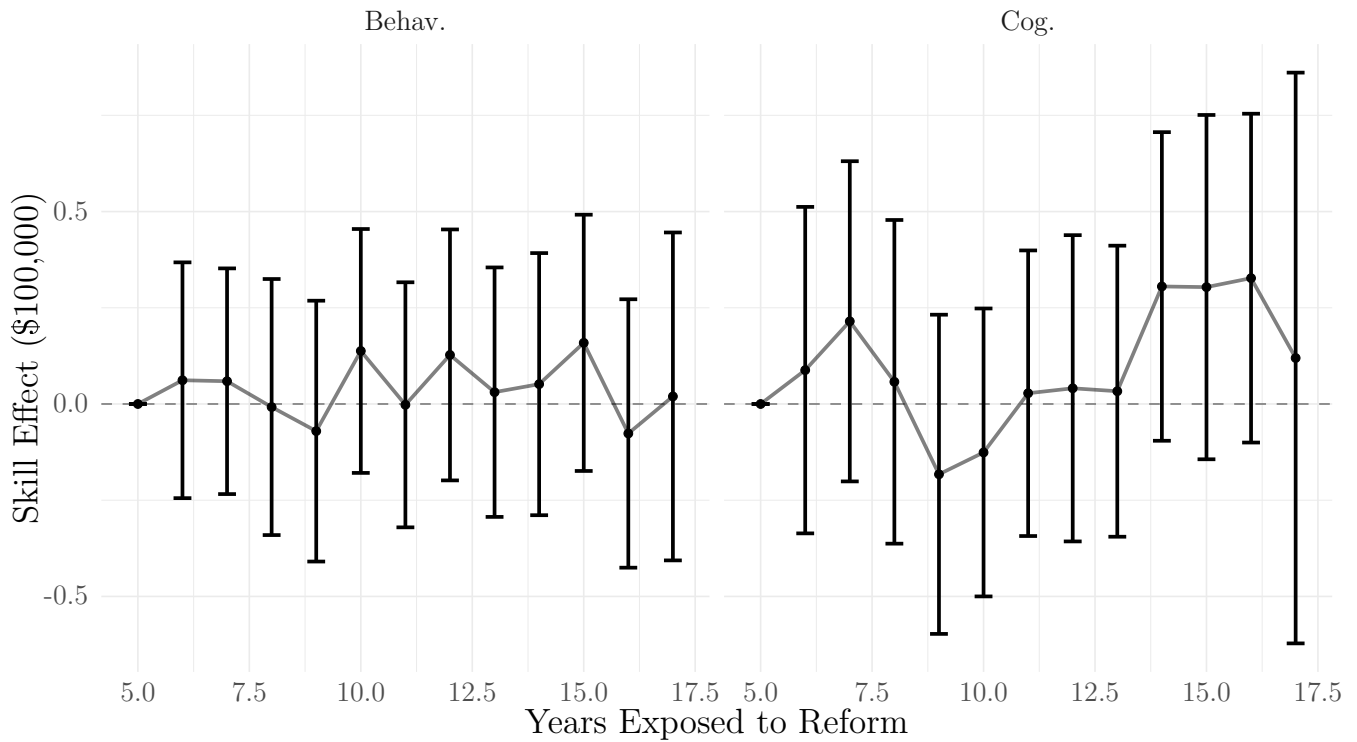
This figure plots estimates of the effect of exposure to welfare reform on cognitive and behavioral skills in the PSID-CDS data using a sibling design. The regression specification is: $\hat{\theta}_{j,m,f,t,1} = \mu_m + \sum_e \beta_e \mathbf{1}\{\text{Exposure}_{m,f,t} = e\} + \epsilon_{j,m,f,t}$ where μ_m is a mother fixed effect and $\text{Exposure}_{m,f,t}$ is the number of years between ages 0 and 17 that the child experienced a policy environment with time limits. Error bars show 95% confidence intervals with standard errors clustered at the level of mothers.

social insurance. The results challenge researchers and policymakers to think about child skill development as a first order component when weighing the costs and benefits of social safety net policies. The relevance of both money and time for skill development has unique – and in this case quite dramatic – implications for the design of work incentives, so it is especially important to have robust and externally valid estimates of these parameters. While this paper takes care to produce credible estimates that are internally consistent with the data and sample of interest, more work is required to reconcile evidence across studies and populations. For example, Mullins (2020) uses experimental assignment to welfare-to-work programs and finds no evidence in that sample that non-maternal care has negative impacts on children. Furthermore, while a *ceterus paribus* effect of employment is likely suitable for studying policies that only shift employment decisions, it does not allow for the study of policies that affect childcare arrangements. The three most important directions for future research are therefore to (1) expand the study to a broader population of interest and model selection into and out of these populations through endogenous household formation and fertility; (2) allow for childcare arrangements to be endogenous and consider policies that may adjust decisions along this margin (for example childcare subsidies); and (3) to reconcile conflicting evidence across studies on the role of maternal employment on child outcomes to produce more robust evidence for policy analysis.

References

AIZER, A., S. ELI, J. FERRIE, AND A. LLERAS-MUNEY (2016): "The Long-Run Impact of Cash Transfers to Poor Families," *American Economic Review*, 106, 935–71.

AIZER, A., H. W. HOYNES, AND A. LLERAS-MUNEY (2022): "Children and the US Social Safety Net: Balancing Disincentives for Adults and Benefits for Children," Working Paper 29754, National Bureau of Economic Research.

AKEE, R., W. COPELAND, E. J. COSTELLO, AND E. SIMEONOVA (2018): "How Does Household Income Affect Child Personality Traits and Behaviors?" *American Economic Review*, 108, 775–827.

ARCIDIACONO, P. AND R. A. MILLER (2020): "Identifying dynamic discrete choice models off short panels," *Journal of Econometrics*, 215, 473–485.

ATKINSON, A. B. AND J. E. STIGLITZ (1976): "The design of tax structure: direct versus indirect taxation," *Journal of Public Economics*, 6, 55–75.

ATTANASIO, O., P. LEVELL, H. LOW, AND V. SÁNCHEZ-MARCOS (2018): "Aggregating Elasticities: Intensive and Extensive Margins of Women's Labor Supply," *Econometrica*, 86, 2049–2082.

ATTANASIO, O., C. MEGHIR, AND E. NIX (2020): "Human Capital Development and Parental Investment in India," *The Review of Economic Studies*, 87, 2511–2541.

BAILEY, M. J., H. W. HOYNES, M. ROSSIN-SLATER, AND R. WALKER (2020): "Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program," Working Paper 26942, National Bureau of Economic Research.

BAILEY, M. J., S. SUN, AND B. TIMPE (2021): "Prep School for Poor Kids: The Long-Run Impacts of Head Start on Human Capital and Economic Self-Sufficiency," *American Economic Review*, 111, 3963–4001.

BARR, A., J. EGGLESTON, AND A. A. SMITH (2022): "Investing in Infants: the Lasting Effects of Cash Transfers to New Families," *The Quarterly Journal of Economics*, qjac023.

BASTIAN, J. AND L. LOCHNER (2021): "The EITC and Maternal Time Use: More Time Working and Less Time with Kids?" *Journal of Labor Economics*, 0, null.

BASTIAN, J. AND K. MICHELMORE (2018): "The Long-Term Impact of the Earned Income Tax Credit on Children's Education and Employment Outcomes," *Journal of Labor Economics*, 36, 1127–1163.

BECKER, G. AND H. G. LEWIS (1973): "On the Interaction between the Quantity and Quality of Children," *Journal of Political Economy*, 81.

BECKER, G. AND N. TOMES (1976): "Child Endowments and the Quantity and Quality of Children," *Journal of Political Economy*, 84.

BENABOU, R. (2002): "Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?" *Econometrica*, 70, 481–517.

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

BERNAL, R. AND M. P. KEANE (2010): "Quasi-structural estimation of a model of childcare choices and child cognitive ability production," *Journal of Econometrics*, 156, 164–189.

——— (2011): "Child Care Choices and Children's Cognitive Achievement : The Case of Single Mothers," *Journal of Labor Economics*, 29, 459–512.

BHANDARI, A., D. EVANS, M. GOLOSOV, AND T. SARGENT (2021): "Efficiency, Insurance, and Redistribution Effects of Government Policies," Tech. rep., working paper.

BITLER, M. P., J. B. GELBACH, H. W. HOYNES, AND M. ZAVODNY (2004): "The impact of welfare reform on marriage and divorce," *Demography*, 41, 213–236.

BLUNDELL, R., M. COSTA DIAS, C. MEGHIR, AND J. SHAW (2016): "Female Labor Supply, Human Capital, and Welfare Reform," *Econometrica*, 84, 1705–1753.

BLUNDELL, R. AND A. SHEPHARD (2011): "Employment, Hours of Work and the Optimal Taxation of Low-Income Families," *The Review of Economic Studies*, 79, 481–510.

BONHOMME, S., K. JOCHMANS, AND J.-M. ROBIN (2016): "Estimating multivariate latent-structure models," *The Annals of Statistics*, 44, 540–563.

——— (2017): "Nonparametric estimation of non-exchangeable latent-variable models," *Journal of Econometrics*, 201, 237–248.

BONHOMME, S. AND E. MANRESA (2015): "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 83.

BRILLI, Y. (2022): "Mother's Time Allocation, Childcare, and Child Cognitive Development," *Journal of Human Capital*, 16, 233–272.

BRUINS, M. (2019): "Taxes, Welfare and the Resources Parents Allocate to Children," *Working Paper*.

CAUCUTT, E. M. AND L. LOCHNER (2020): "Early and late human capital investments, borrowing constraints, and the family," *Journal of Political Economy*, 128, 1065–1147.

CAUCUTT, E. M., L. LOCHNER, J. MULLINS, AND Y. PARK (2020): "Child Skill Production: Accounting for Parental and Market-Based Time and Goods Investments," Working Paper 27838, National Bureau of Economic Research.

CHAN, M. K. (2013): "A Dynamic Model of Welfare Reform," *Econometrica*, 81, 941–1001.

CHERNOZHUKOV, V. AND H. HONG (2003): "An MCMC approach to classical estimation," *Journal of Econometrics*, 115, 293–346.

CHETTY, R., J. N. FRIEDMAN, N. HILGER, E. SAEZ, D. W. SCHANZENBACH, AND D. YAGAN (2011): "How does your kindergarten classroom affect your earnings? Evidence from Project STAR," *The Quarterly journal of economics*, 126, 1593–1660.

CHETTY, R., J. N. FRIEDMAN, AND E. SAEZ (2013): "Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings," *American ECconomic Review*, 103.

CUNHA, F., J. HECKMAN, AND S. SCHENNACH (2010): "Estimating the Technology of Cognitive and Noncognitive Skill Formation." *Econometrica*, 78, 883–931.

CUNHA, F. AND J. J. HECKMAN (2008): "Formulating , Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Journal of Human Resources*, 43.

DAHL, G. B. AND L. LOCHNER (2012): "The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit," *American Economic Review*, 102, 1927–1956.

DEL BOCA, D., C. FLINN, AND M. WISWALL (2014): "Household Choices and Child Development," *Review of Economic Studies*, 81.

DIAMOND, P. (1980): "Income taxation with fixed hours of work," *Journal of Public Economics*, 13, 101–110.

DUNCAN, G. J., P. A. MORRIS, AND C. RODRIGUES (2011): "Does money really matter? Estimating impacts of family income on young children's achievement with data from random-assignment experiments." *Developmental psychology*, 47, 1263–79.

FEENBERG, D. R. AND E. COUTTS (1993): "An Introduction to the TAXSIM Model," *Journal of Policy Analysis and Management*, 12, 189–194.

FLODEN, M. (2001): "The effectiveness of government debt and transfers as insurance," *Journal of Monetary Economics*, 48, 81–108.

FLOOD, S., M. KING, R. RODGERS, S. RUGGLES, J. R. WARREN, AND M. WESTBERRY (2021): "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]," .

GARCÍA, J. L., J. J. HECKMAN, D. E. LEAF, AND M. J. PRADOS (2020): "Quantifying the Life-Cycle Benefits of an Influential Early-Childhood Program," *Journal of Political Economy*, 128, 2502–2541.

GAYLE, G.-L., L. GOLAN, AND M. SOYTAS (2015): "What Accounts for the Racial Gap in Time Allocation and Intergenerational Transmission of Human Capital?" .

GENNETIAN, L. A. AND V. KNOX (2003): "Staying Single: The Effects of Welfare Reform Policies on Marriage and Cohabitation," Tech. rep., Manpower Demonstration Research Corporation.

GRIFFEN, A. S. (2019): "Evaluating the effects of childcare policies on children's cognitive development and maternal labor supply," *Journal of Human Resources*, 54, 604–655.

GROGGER, J. (2002): "The Behavioral Effects of Welfare Time Limits," *American Economic Review*, 92.

——— (2003): "The Effects of Time Limits, the EITC, and Other Policy Changes on Welfare Use, Work, and Income among Female-Headed Families," *Review of Economics and Statistics*, 85, 394–408.

GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2020): "Child-related transfers, household labour supply, and welfare," *The Review of Economic Studies*, 87, 2290–2321.

HEATHCOTE, J., K. STORESLETTEN, AND G. VIOLANTE (2017): "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 132.

HECKMAN, J., R. PINTO, AND P. SAVELYEV (2013): "Understanding the mechanisms through which an influential early childhood program boosted adult outcomes," *American Economic Review*, 103, 2052–86.

HECKMAN, J. J., S. HYEOK, R. PINTO, A. PETER, S. H. MOON, P. A. SAVELYEV, AND A. YAVITZ (2010): "The rate of return to the HighScope Perry Preschool Program," *Journal of Public Economics*, 94, 114–128.

HECKMAN, J. J., J. STIXRUD, AND S. URZUA (2006): "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior," *Journal of Labor Economics*.

HO, C. AND N. PAVONI (2020): "Efficient Child Care Subsidies," *American Economic Review*, 110, 162–99.

HOYNES, H. (1996): "Welfare Transfers in Two-Parent Families: Labor Supply and Welfare Participation under AFDC-UP," *Econometrica*, 64.

HOYNES, H., D. W. SCHANZENBACH, AND D. ALMOND (2016): "Long-Run Impacts of Childhood Access to the Safety Net," *American Economic Review*, 106, 903–34.

HSIN, A. AND C. FELFE (2014): "When does time matter? Maternal employment, children's time with parents, and child development," *Demography*, 51, 1867–1894.

HU, Y. AND M. SHUM (2012): "Nonparametric identification of dynamic models with unobserved state variables," *Journal of Econometrics*, 171, 32–44.

KASAHARA, H. AND K. SHIMOTSU (2009): "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, 77, 135–175.

KEANE, M. P. AND K. I. WOLPIN (2010): "The role of labor and marriage markets, preference heterogeneity, and the welfare system in the life cycle decisions of black, hispanic, and white women," *International Economic Review*, 51, 851–892.

KEARNEY, M. S. (2004): "Is There an Effect of Incremental Welfare Benefits on Fertility Behavior? A Look at the Family Cap," *The Journal of Human Resources*, 39, 295–325.

KLINE, P. AND C. R. WALTERS (2016): "Evaluating public programs with close substitutes: The case of Head Start," *The Quarterly Journal of Economics*, 131, 1795–1848.

LOW, H., C. MEGHIR, L. PISTAFERRI, AND A. VOENA (2018): "Marriage, Labor Supply and the Dynamics of the Social Safety Net," Working Paper 24356, National Bureau of Economic Research.

MARSCHAK, J. (1953): *Economic Measurements for Policy and Prediction*, John Wiley & Sons.

MCFADDEN, D. (1981): "Econometric models of probabilistic choice," *Structural analysis of discrete data with econometric applications*, 198272.

MEYER, B. D. (2002): "Labor Supply at the Extensive and Intensive Margins : The EITC , Welfare , and Hours Worked," *American economic review*, 92.

MIRRLEES, J. A. (1971): "An Exploration in the Theory of Optimum Income Taxation," *The Review of Economic Studies*, 38, 175–208.

MOFFITT, R. A. (2015): "The Deserving Poor, the Family, and the U.S. Welfare System," *Demography*, 52, 729–749.

MULLINS, J. (2020): "A Structural Meta-Analysis of Welfare to Work Programs and their Impacts on Children," *Working Paper*.

PATTERSON, C. (2022): "The Matching Multiplier and the Amplification of Recessions," *Working Paper*.

SAEZ, E. (2001): "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 68, 205–29.

——— (2002): "Optimal Income Transfer Programs: Intensive Versus Extensive Labor Supply Responses," *Quarterly Journal of Economics*, 117, 1039–1073.

SAEZ, E. AND S. STANTCHEVA (2016): "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 106, 24–45.

VAN DIJK, W. (2019): "The socio-economic consequences of housing assistance," *Working Paper*.

A Figures

Figure 13: Model Selection Information Criteria

This figure shows the Akaike and Bayesian Information Criteria (AIC & BIC) for different choices of the number of types (K) in the model. The formulae for these model selection criteria are $AIC = 2p - L_M(\hat{\theta})$ and $AIC = p\ln(M) - 2L_M(\hat{\theta})$ where $p = \dim(\theta)$ and M is the number of observations.

Figure 14: Prior and Posterior for Cobb-Douglas Shares

This figure plots the posterior over Cobb-Douglas shares from the Quasi-Bayesian estimation as compared to pre-specified priors. A Metropolis-Hastings algorithm is used to sample from the posterior distribution. See Appendix D.5 for more details.

Figure 15: Model Fit by Latent Type for Log Monthly Earnings

This figure shows the model's fit of mean rates of log monthly earnings by latent type. Means from the data are calculated using posterior weights generated by the model's estimates.

Figure 16: Average Observable Characteristics By Type

This figure plots average observable characteristics for each type k. "AFDC" is the mean rate of participation in cash welfare across all individual-year observations. "Age at First Birth" is the mother's age at the birth of her first child. "College" and "High School" are the fraction of individuals in each group who have at least 16 and 12 years of education respectively in 1997. "LFP" is the mean rate of labor force participation across all individual-year pairs. "Total Births" is the average number of total births for each mother as of the year 2010. Error bars show 95% confidence intervals.

Figure 17: Selection into Program Participation by Latent Characteristics

This figure plots average program participation for each type k against three type-specific latent characteristics: (1) costs of work ($\alpha_{H,k}$); (2) log of the cost of program participation ($\log(\alpha_{A,k})$); and (3) mean monthly wages at age 25 ($\gamma_{W,0,k} + \gamma_{W,1,k}25$). Point sizes are proportional to estimated population probabilities for each type.

Figure 18: Effect of \$1000 in Annual Household Income on Cognitive and Behavioral Skills

This figure shows the average effect of an unconditional \$1000 transfer implied by production parameter estimates. Effect sizes are expressed as fractions of a standard deviation of the cognitive and behavioral score. The "Ceterus Paribus" effect is the average effect on skills (in fractions of a standard deviation) of an additional \$1000 in household income, fixing the choice probabilities of the model. The "Endogenous" effect is the average effect on skills when choice probabilities adjust to the \$1000 unconditional transfer. Calculations are made for the oldest CDS child in each household. For behavioral skills, the estimate from Akee et al. (2018) is constructed by combining the total effect in Table 3 with the average effect on income taken from Table 2, normalized to a \$1000 transfer. The estimate from Dahl and Lochner (2012) is taken from Table 3 of that paper. To account for time aggregation, effect sizes for cited papers are divided by two for both studies and 95% confidence intervals are calculated using reported standard errors. For model estimated effects, 95% credibility intervals are shown, and are calculated by bootstrap sampling from the MCMC chain generated in Section 4.4.

Figure 19: Effect of Work on Cognitive Skills, Children Aged 0-5

This figure shows the estimated ceterus paribus effects of work on cognitive skills, expressed as fraction of a standard deviation. This effect is given by $\delta_{\tau,a,C} \log\left(\frac{112-30}{112}\right)$ for each age a. Point estimates and standard deviations for Bernal and Keane (2011) are from column (6) of Table 6 and Table 3 in that paper. For model estimates, 95% credibility intervals are shown, and are calculated by bootstrap sampling from the MCMC chain generated in Section 4.4.

Figure 20: The Total Developmental Effect of Employment

This figure replicates Figure 6 using both the Quasi-Bayesian estimates of production parameters and the GMM estimates that impose only the strict IV moment conditions. 95% credibility intervals for the preferred estimates are calculated as in Figure 6, while 95% confidence intervals for the GMM estimates are calculated by simulating draws from the estimated asymptotic distribution.

Figure 21: Cobb-Douglas Share Estimates

This figure compares the GMM estimates of Cobb-Douglas shares for time and money inputs in the baseline specification ($K = 10$) to the specification with simplied heterogeneity ($K = 2$). "IV - Model" refers to the GMM estimates using moment conditions described in (50), "IV-Strict" refers to estimates from the moment conditions using only policy variation as in (51).

Figure 22: Effect of "No Welfare Reform" counterfactual on Work and Program Participation

Figure 23: Heterogeneous Effects of "No Welfare Reform" on Skill Outcomes
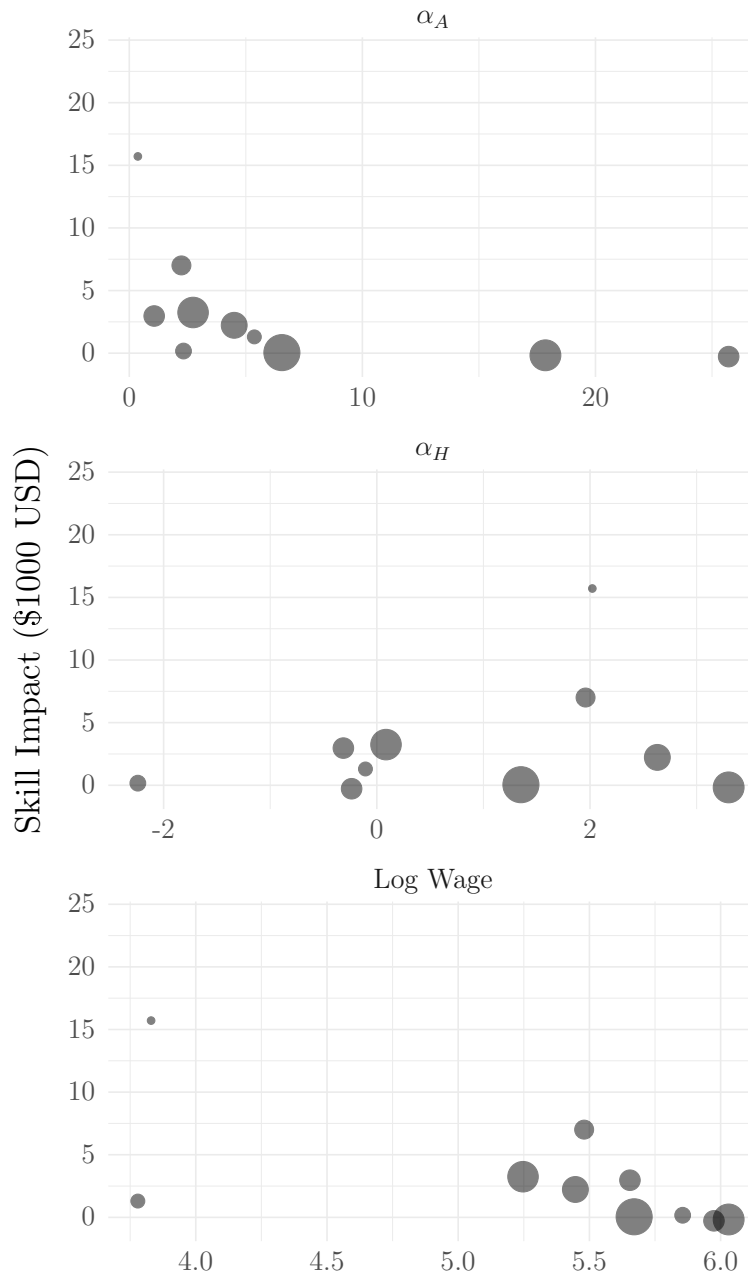
This figure plots the total skill impacts (in thousands of year 2000 USD) for children of each type resulting from the "No Welfare Reform" policy counterfactual (in which the policy environment is paused in 1996) against each type's cost of program participation (α_A), cost of work (α_H), and mean log wages at age 25 ("Log Wage").

B   Description of Transfer Policies

This section describes the computation of the transfer functions (T^F, T^A, T^T). The details of this section very closely follow Chan (2013), which should be consulted for further details.

Welfare

The transfer function T^A includes a benefit computation, and an eligibility test:

$$T^A_{mt} = \text{El}^A_{mt} \times \text{Ben}^A_{mt} \tag{33}$$

Where

$$\text{El}^A_{mt} = \mathbf{1}\{\omega_{mt} \le \Omega_{mt}\} \times \mathbf{1}\left\{E_{mt} + N_{mt} < r_{A,g,mt}e_{A,mt}\right\} \times \mathbf{1}\left\{(E_{mt} - D_{Ae,mt})(1 - R_{Ae,mt}) + N_{mt} < r_{An,mt}e_{A,mt}\right\}. \tag{34}$$

Eligibility above is defined as the combination of a time limit, a net income test, and a gross income test. Both tests compare income with a need standard, $e_{A,mt}$ which is inflated by some rate $(r_{Ag,mt}, r_{Ae,mt})$. Second, the computation of net income involves a fixed disregard on earnings, $D_{Ae,mt}$ and a percentage disregard. Benefit computation follows similarly:

$$\text{Ben}^A_{mt} = \max\{G_{A,mt} - (E_{mt} - D_{Ab,mt})(1 - R_{Ab,mt}) - N_{mt}, 0\} \tag{35}$$

The payment standard $G_{A,mt}$ sets the generosity of the program when no other sources of income are reported, while the dollar and percentage disregards $(D_{Ab,mt}, R_{Ab,mt})$ combine to determine net income. Importantly, these policy parameters are a function of the mother's state of residence, the number of dependant children and the year. In this model, these variables are all a function of the mother-year index, mt. In reality, individuals may be subject to asset tests for eligibility, which are not modelled here. Time limits in some states are also *periodic* in the sense that in addition to the total limit, there are shorter limits on the number of months of consecutive use. These periodic time limits are also not modelled due to tractability concerns.

Food Stamps

Similarly to welfare, the food stamp transfer function T^F can be written as:

$$T^A_{mt} = \text{El}^F_{mt} \times \text{Ben}^F_{mt} \tag{36}$$

Where

$$\text{El}^F_{mt} = \mathbf{1}\left\{E_{mt} + N_{mt} < 1.3e_{F,mt}\right\} \times \mathbf{1}\{\underbrace{0.8E_{mt} + N_{mt} + \text{Ben}^A_{mt} - 134}_{=\text{Net}^F_{mt}} < e_{F,mt}\}. \tag{37}$$

In the above expression, $e_{F,mt}$ is referred to as the poverty guideline, and the net income includes a standard 20% disregard and \$134 deduction. While the true food stamp benefit formula technically allows for further deductions for child care expenses, child support payments, and shelter expenses, I have insufficient data to calculate these deductions. Finally, given a maximum benefit $G_{F,mt}$, the benefit calculation is:

$$\text{Ben}^F_{mt} = \max\left\{G_{F,mt} - 0.3\text{Net}^F_{mt}, 0\right\}. \tag{38}$$

Data Sources for Program Rules

To summarize, the parameter vector Z_{mt}^A can be written as

$$Z_{mt} = \{r_{Ag,mt}, r_{Ae,mt}, e_{A,mt}, D_{Ae,mt}, R_{Ae,mt}, G_{A,mt}, D_{Ab,mt}, R_{Ab,mt}, \mathcal{L}_{mt}\},$$

while Z_{mt}^F can be summarized as

$$Z_{mt}^F = \{e_{F,mt}, G_{F,mt}\}.$$

Parameters on welfare that comprise Z_{mt}^A and Z_{mt}^F were collected from the Urban Institute's TRIM3 simulation database[19] for years 1985-2011. In addition, since rules on net income calculations were much more simple prior to 1993, I use a 30% disregard across all states [20]. Mothers were merged with program rules based on their state of residence, the year, and the number of children in their household of age 17 or younger.

Taxes

Taxes consist of a federal and a state computation. When earned income is sufficiently low, T^T will arrive in the form of a net payment (when income tax obligations are exceeded by the EITC). In theory, the relevant parameters to compute taxes include those that define the federal and state EITC programs, state and federal deductions and exemptions, and the marginal income tax rate with their corresponding brackets for state and federal income tax. In practice, I use the TAXSIM model of Feenberg and Coutts (1993), to approximate the tax function. Given the relevant year, state, and family size (in our model these are all exogenous functions of the index, mt), TAXSIM computes $T_{mt}^T(e)$ for any given earnings level. Thus, for each mt in my sample, I compute $T_{mt}^T(e)$ for earnings levels e on a grid, using increments of \$100, between \$0 and \$100,000[21]. Using this grid, the tax function is approximated using linear interpolation between these grid points.

C   Model Solution with General Technology

Here consider the solution of any dynamic model with a general state S, a general action c, and a generalization of the technology derived in the main section:

$$\theta_{f,t+1,j} = \exp(\mu_{\theta,a,j} + \eta_{f,t,j}) \prod_{l=1}^{L_1} \tau_{f,t,l}^{\delta_{\tau,a(f,t),j,l}} \prod_{l=1}^{L_2} \tilde{\tau}_{t,l}^{\delta_{\tilde{\tau},a(f,t),j,l}} x_{f,t}^{\delta_{x,a(f,t),j}} \theta_{f,t,C}^{\delta_{\theta,C,j}} \theta_{f,t,B}^{\delta_{\theta,B,j}}$$

---

[19]Source: http://trim3.urban.org/

[20]This approach is taken also in Chan (2013)

[21]This suits as a reasonable upper bound in my sample

where there are now L_1 private time investment categories and L_2 public time investment categories. The constraints on investment can be written as:

$$x \le \mathcal{C}_x(c) \tag{39}$$

$$\sum_l^{L_1} \sum_f \tau_{f,l} + \sum_l^{L_2} \tilde{\tau}_l \le \mathcal{C}_\tau(c) \tag{40}$$

One can easily show that if the value at $t+1$ can be written as:

$$V_{t+1}(S,\theta) = \nu_{t+1}(S) + \sum_f \sum_{j \in \{C,B\}} \alpha_{V,a(f,t)+1,j} \log(\theta_{f,t+1,j})$$

then the solution at t can be written as[22]:

$$V_t(S,\theta) = \max_{[c,\tau,x] \in \mathcal{C}(S)} \left\{ U_t(S,c) + \beta \mathbb{E}_{S'|S,c}[\nu_{t+1}(S')] + \beta \sum_j \sum_f \alpha_{V,a(f,t)+1,j} \Big(\delta_{x,a(f,t),j} \log(x_f) \right.$$

$$\left. + \sum_{l=1}^{L_1} \delta_{\tau,a(f,t),j,l} \log(\tau_{f,l}) + \sum_{l=1}^{L_2} \delta_{\tilde{\tau},a(f,t),l} \log(\tilde{\tau}_l) \Big) \right\}$$

$$+ \sum_{j \in \{C,B\}} \sum_f (\alpha_{\theta,j} + \beta \delta_{\theta,j,C} \alpha_{V,a(f,t)+1,C} + \beta \delta_{\theta,j,B} \alpha_{V,a(f,t)+1,B}) \log(\theta_{f,t,j})$$

$$+ \sum_{j \in \{B,C\}} \sum_f \beta \alpha_{V,a(f,t)+1,j} (\mu_{\theta,a(f,t),j} + \mathbb{E}[\eta_{f,t,j}]). \tag{41}$$

As long as the evolution of the state S is not affected by θ, then additive separability carries through. The first two lines of this equation give the formula for $\nu_t(S)$, the third line gives the recursive formula for each $\alpha_{V,a,j}$, and the final line gives the formula for the additive, invariant constant C_t.

Further notice that the solution for optimal time investment is given by:

$$\tau_{f,t,l} = \phi_{\tau,f,t,l} \mathcal{C}_\tau(c), \qquad \tilde{\tau}_{t,l} = \phi_{\tilde{\tau},t,l} \mathcal{C}_\tau(c)$$

where

$$\phi_{\tau,f,t,l} = \frac{\beta \sum_j \alpha_{V,a(f,t)+1,j} \delta_{\tau,a(f,t),j,l}}{\sum_j \sum_f \beta \alpha_{V,a(f,t)+1,j} \left(\sum_{l=1}^{L_1} \sum_f \delta_{\tau,a(f,t),j,l} + \sum_{l=1}^{L_2} \delta_{\tilde{\tau},a(f,t),j,l} \right)} \tag{42}$$

$$\phi_{\tilde{\tau},t,l} = \frac{\beta \sum_f \sum_j \alpha_{V,a(f,t)+1,j} \delta_{\tilde{\tau},a(f,t),j,l}}{\sum_j \sum_f \beta \alpha_{V,a(f,t)+1,j} \left(\sum_{l=1}^{L_1} \sum_f \delta_{\tau,a(f,t),j,l} + \sum_{l=1}^{L_2} \delta_{\tilde{\tau},a(f,t),j,l} \right)} \tag{43}$$

$$\tag{44}$$

Since similar logic applies for money investment, the outcome equation for the vector of skills can be written as:

$$\log(\theta_{f,t+1}) = \mu_{\theta,a} + \delta_{x,a} \log(\mathcal{C}_x(c_t)) + \bar{\delta}_\tau \log(\mathcal{C}_\tau(c_t)) + e(a,\boldsymbol{a}) + \eta_{f,t}$$

---

[22]One can set $\alpha_\theta = 1$ without loss of generality since it is a free scaling parameter.

where

$$\overline{\delta}_{\tau,a} = \left(\sum_{l=1}^{L_1} \delta_{\tau,a,l} + \sum_{l=1}^{L_2} \delta_{\tilde{\tau},a,l} \right)$$

is the sum of the factor shares of all time investment categories. Thus, when forecasting the welfare or skill outcome impacts of policies that affect incentives only through adjustments to the constraint set $\mathcal{C}$, only the sum of the factor shares is relevant. In this sense, the assumptions made in the main body of the paper with respect to technology and investment categories are without loss of generality within the Cobb-Douglas class, and adhere to Marschak's maxim (Marschak, 1953).

D   Identification and Estimation

D.1   Identification and Estimation of Skill Measurement Equations

This paper follows the approach described by Cunha and Heckman (2008) to appropriately handle measurement error when estimating the key technology parameters, $(\delta_x, \delta_\tau, \delta_\theta)$. Since the factors have no fixed location or scale, the normalization $\mathbb{E}[\theta_C] = \mathbb{E}[\theta_B] = 0$ applies alongside the normalization that $\lambda_{C,1} = \lambda_{B,1} = 1$. With these in hand, the measurement assumptions yield the following set of identifying relationships[23]:

$$\mu_{j,l} = \mathbb{E}[\tilde{\theta}_{j,l}], \qquad j \in \{C, B\}, \ l \in \{1, 2\} \tag{45}$$

$$\lambda_{j,2} = \frac{\mathbb{C}(\tilde{\theta}_{j,1997,2}, \tilde{\theta}_{j,2002,1})}{\mathbb{C}(\tilde{\theta}_{j,1997,1}, \tilde{\theta}_{j,2002,1})}, \qquad j \in \{C, B\} \tag{46}$$

The sample analog of these relationships (replacing population moments with sample means, variances, and covariances) yield consistent estimates of the measurement parameters, $(\hat{\mu}, \hat{\lambda})$, which in turn allows the construction of two factor scores for each measurement:

$$\hat{\theta}_{j,m,f,t,1} = \tilde{\theta}_{j,m,f,t,1} - \hat{\mu}_{j,1}, \qquad \hat{\theta}_{j,m,f,t,2} = \frac{\tilde{\theta}_{j,m,f,t,2} - \hat{\mu}_{j,2}}{\hat{\lambda}_{j,2}}.$$

D.2   Anchoring Skills to Earnings and Crime Outcomes

This section describes the formal assumptions and analysis that lead to the choice of anchoring coefficients described in Section 4.1.

D.2.1   A Useful Benchmark

Heckman et al.'s (2013) mediation analysis of the effects of Perry on adult outcomes provides a useful source of validating evidence for the proceeding anchoring exercises, in that it identifies the

---

[23]The factor loadings λ are not uniquely identified by these particular covariance restrictions, but applying one restriction per factor loading provides the simplest estimator.

Table 7: Value of a Standard Deviation Increase in Behavioral Skills: Perry Preschool Program

	Males		Females
Effect Size on Behavioral Skills[(a)]	0.21		0.29
	Earnings	Crime	Crime
Effect Size on Outcome[(b)]	0.60	-0.32	-0.32
% Effect on Outcome[(b)]	61.2	-48.1	-100
Value of total treatment effect[(c)]	$36,705	$87,074	$14,663
% Treatment effect attributable to skill[(d)]	20%	30%	60%
% Effect on Outcome of 1 s.d. increase in Behavioral Skill[(e)]	58.6	-68.9	-202.7
Value of 1 s.d. increase in behavioral skill[(f)]	$35,105	$124,919	$29,716

Notes: (a) "Effect Size" is the fraction of a standard deviation effect on the skill outcome. Effect size Δ_X for outcome X are derived from one-sided p-values reported in Figure 5 of Heckman et al. (2013) using the formula $1 - p = \Phi(\sqrt{N}\Delta_X)$. (b) Source: Table (8) of Heckman et al. (2010), using Number of Felony Arrests as index of crime. (c) Source: Table (8) Heckman et al. (2010), deflated to year 2000 dollars. (d) Figures (6) and (7) of Heckman et al. (2013) suggest behavioral skills explain between 20 and 70% of reductions in male criminal behavior, about 20% of increases in male earnings, and 60-70% of reductions in female criminal behavior. (e) Formula is Fraction Attributable $\times$ % Effect $\div$ Effect Size on Skills (f) Formula is Fraction Attributable $\times$ Value of Treatment Effect (year 2000 $USD) $\div$ Effect Size on Skills where lower bounds are chosen on fractions attributable.

net present value returns to a standard deviation in behavioral skills for the domains of earnings and criminal behavior. In the latter case, those findings will be combined explicitly with available evidence on criminal behavior in the CDS sample to form a net present value calculation. Table 7 documents the results of this study and derives a net present dollar valuation for one standard deviation of behavioral skills. The table focuses on behavioral skills because Heckman et al. (2013) do not find a substantial role for cognitive skills in explaining the effects of Perry. Nevertheless, the results provide a useful guide for benchmarking and choosing the total resource value of skills.

D.2.2 Lifetime Earnings

The proportionality of skill returns on earnings at each age implies the following model for the net present value of earnings:

$$\mathcal{Y}^{EARN}(\theta) = (1 + \tilde{\gamma}_{E,C} \log(\theta_C) + \tilde{\gamma}_{E,B} \log(\theta_B)) \times \sum_{a=18}^{65} \beta^{a-18} \gamma_{Y,a}$$

To forecast the age effects $\gamma_{Y,a}$ this paper uses the 1997-2018 waves of the Annual Social and Economic Supplement of the Current Population Survey (Flood et al., 2021) to estimate a model

Figure 24: Lifecycle Earnings Estimates by Cohort

This figure shows average earnings by age for each of the 5 cohorts from the 1997-2018 waves of ASEC. It also shows the resulting forecast of average earnings for the cohort born after 1985 when using estimates from the quadratic model described in the text.

with cohort effects and a quadratic age trend:

$$Y_i = \kappa_c + \tilde{\gamma}_0 a_i + \tilde{\gamma}_1 a_i^2 + \epsilon_i$$

where κ_c are cohort effects for birth cohorts in the set: $\{\leq 1950, 1951 - 1960, 1961 - 1970, 1971 - 1980, \geq 1985\}$. This model provides a forecast of average earnings for the 1985-1997 cohort. Figure 24 shows estimates of this projection as well as the age profile of earnings for the different cohorts in the sample.

The earnings model also identifies the $\gamma_{Y,a}$ up to a constant of proportionality when using the 1985-1997 cohort effect:

$$\hat{\gamma}_{Y,a} = \frac{\mu_{CDS}}{\mu_{USA}} (\kappa_{\geq 1980} + \tilde{\gamma}_0 a + \tilde{\gamma}_1 a^2)$$

where the the constant of proportionality μ_{CDS}/μ_{USA} reflects the fact that average skills in the selected CDS sample is different from the representative ASEC sample. A comparison of mean earnings between the ages of 23 and 27 (for which observations of the CDS sample are available) provides an estimate of this ratio. Table 8 shows the results and demonstrates the resulting

Table 8: Calculation of NPV of Lifecycle Earnings for CDS Sample

	Sample	
Age	PSID-CDS	CPS (1985-)
23	6736	11971
24	7732	14223
25	7680	17651
26	9689	18640
27	13036	20355
Mean (23-27)	8975	16568

$\widehat{\mu_{CDS}/\mu_{USA}} \approx 0.54$

$\widehat{NPV}_{USA} = \$723,298$

$\widehat{NPV}_{CDS} = \$391,803$

This table reports mean earnings for individuals aged 23-27 from the selected PSID-CDS sample used in this paper and from the 1997-2018 waves of the ASEC. Calculations for the latter case include only individuals born 1985 or after. The lower panel shows calculations of the net present value of earnings for the CDS sample by taking the estimated NPV from the ASEC samples and multiplying by the ratio of average earnings in the two samples (see text for more details).

calculation of the ratio μ_{CDS}/μ_{USA} used to calculate the net present value of earnings for CDS sample members.

This approach delivers an estimated net present value of lifetime earnings for the ASEC sample of \$723,298 in year 2000 dollars, using a chosen value of β of 0.98. [24] Based on observed earnings for the CDS estimation sample, a value of just over half this number (\$391,800) is chosen for anchoring the earnings returns to skills.

With this number in hand we can estimate the effect of cognitive and behavioral skills on lifetime earnings by estimating their proportional effect on earnings in early adulthood. Using the most recent available measure of skills for each child, I regress annual earnings on the first factor score for cognitive and non-cognitive ability, using the second measure of each ability as instruments

---

[24]In Chetty et al. (2011) the authors use a discount rate of 3%, finding a net present value of \$522,000 in 2010 dollars when discounted to age 12. Making an equivalent calculation using my method delivers a comparable present value of \$599,000.

to account for measurement error. Table 9 shows the results using all available observations of CDS children (column 1), as well as using only children from the chosen sample of single mothers (columns 2 & 3). Coefficients on cognitive and non-cognitive factors can be interpreted as the effect of a one standard deviation increase in that factor. Table 9 also reports the implied percentage impacts of a standard deviation increase in each skill, which can be interpreted as the pair $(\tilde{\gamma}_{E,C}, \tilde{\gamma}_{E,B})$ and can be combined with estimates of the net present value of earnings to get the anchoring coefficients $(\tilde{\gamma}_{E,C}, \tilde{\gamma}_{E,B})$ reported in Table 2.

Naturally, caution should be applied when attempting to attribute a causal interpretation to these estimates. However, relative to existing results, these appear to be within the range of previous estimates. Using a mediation analysis of the Perry Preschool intervention, Heckman et al. (2013) attribute somewhat larger gains in income to increases in behavioral skills (Table 7). Furthermore, in their analysis of the impacts of Head Start, Kline and Walters (2016) determine that a 10% increase in earnings from a one standard deviation increase in achievement scores is a reasonable estimate. When considering the fact that those changes are evaluated for preschool aged children (as compared to measures of skills much closer to adulthood in this sample) the percentage impacts implied by the estimates in Table 9 appear to be reasonable if not conservative.

D.2.3   Costs of Crime

Similarly to the case of earnings, $\mathcal{Y}^{CRIME}(\theta)$ can be written as:

$$\mathcal{Y}^{CRIME}(\theta) = (1 + \tilde{\gamma}_{CR,C} \log(\theta_C) + \tilde{\gamma}_{CR,B} \log(\theta_B)) \times NPV^{CRIME}$$

While the PSID-CDS does not provide a sufficiently complete account of criminal activity for a direct calculation, it does allow construction of a proxy for the number of arrests for an individual by recording an individual arrest if they report being arrested at least once at any stage in the data, and adding an arrest every time the variable "Age Since Last Arrested" is updated in the data. When using data on the main estimation sample, results in Table 10 suggest that a standard deviation increase in behavioral skills leads to a 16% reduction in the number of arrests. Cognitive skills do not appear to play a role, suggesting that $\gamma_{CR,C} = 0$ is a reasonable choice.

Anchoring is achieved by assuming that this outcome variable is proportional to the overall net present value of criminal behavior, and that this proportional relationship also holds for a comparable measure – the total number of felony arrests – in the Perry Study. Formally, letting NA be the number of arrests:

$$\mathbb{E}[NA(\theta)|CDS] = A(1 + \tilde{\gamma}_{CR,B} \log(\theta_B))$$

$$\mathcal{Y}^{CRIME}_{PERRY}(\theta) = (1 + \tilde{\gamma}_{P,CR,B} \log(\theta_B)) \times NPV^{CRIME}$$
$$\mathbb{E}[NA(\theta)|PERRY] = \tilde{A}(1 + \tilde{\gamma}_{P,CR,C} \log(\theta_C) + \tilde{\gamma}_{P,CR,B} \log(\theta_B))$$

Table 9: Effect of Cognitive and Behavioral Skills on Earnings

	Earnings ($/year)		
	(1)	(2)	(3)
Cognitive	2,907.792***	1,763.632***	1,654.131***
	(427.878)	(582.343)	(607.914)
Behavioral	329.277	839.735*	840.060*
	(415.234)	(441.762)	(441.423)
Age			163.283
			(305.208)
Mother's Ed.			218.983
			(178.856)
Constant	8,364.953***	7,923.075***	1,008.194
	(291.320)	(436.045)	(8,034.083)
Observations	1,644	624	624
	Implied % Effects of 1 s.d. Improvement		
Cognitive	34.7	25.45	23.87
Behavioral	3.93	12.12	12.12
	NPV of 1 s.d. Improvement ($1000 USD year 2000)		
Cognitive	136	99.7	93.5
Behavioral	15.4	47.5	47.5

*p<0.1; **p<0.05; ***p<0.01

This table reports estimates of the relationship between earnings in young adulthood for CDS children and measures of cognitive and behavioral skills. Estimates are calculated by two stage least squares using the first factor score of each skill in the outcome equation and the second factor score as instruments (see section 4.1). Specification (1) shows results for all non-missing observations of CDS children, while Specifications (2) and (3) use only the chosen sample for the paper. The lower panel shows net present value calculations for the value of each skill by taking implied percentage effects on earnings and multiplying by the estimated net present value of earnings in thousands of year 2000 US dollars.

The net present value effects can be scaled then by taking an estimate of the value of one standard deviation of behavioral skills from Heckman et al. (2013), and multiplying by the ratio of estimated percentage effects from the two samples:

$$\gamma_{CR,S} = \underbrace{\tilde{\gamma}_{P,CR,S} NPV^{CRIME}}_{=\$90,000} \times \underbrace{\frac{\tilde{\gamma}_{CR,S}}{\tilde{\gamma}_{P,CR,S}}}_{\approx 0.16/1.35} \, .$$

When averaging over men and women, a one standard deviation increase in behavioral skills in the Perry study is worth \$90,000. Table 7 also suggests that the average reduction in arrests across men and women from a one standard deviation gain in behavioral skills is 135%[25], while in the chosen CDS sample it is 16%. This results in the conservative calculation of \$9,000 (one tenth of the value in Perry).

D.3 Identification of First Stage Parameters

In the first stage the parameters to be identified are those governing the distribution over latent states $(\Pi_K, \Pi_W, \Pi_{W,0})$, those governing the deterministic component of wages, $(\gamma_{W,0}, \gamma_{W,1})$, and those governing preferences $(\alpha_C, \alpha_l, \alpha_\theta, \alpha_H, \alpha_A, \alpha_R, \Gamma_x, \Gamma_\tau)$. Since preference parameters provide a low-dimensional parameterization of choice probabilites, $P(S_{m,t})$, it is sufficient to establish non-parametric identification of those choice probabilities along with the objects that govern the initial distribution of latent states (Π_K and $\Pi_{W,0}$) and transition probabilities (Π_W). In this paper, transition laws for observed state variables are deterministic and known.

While the exact conditions required for identification of these models vary depending on the key sources of variation, a unifying theme is that they require a sufficiently large panel dimension. Hu and Shum (2012) for example show that models of this kind can be identified with a panel of length 5, which these data easily satisfy (see Section 3). The key requirement in Hu and Shum (2012) is a high level condition on the uniqueness of the spectral decomposition of a linear operator defined by the joint distribution of observable and unobservable variables over four periods. Kasahara and Shimotsu (2009) and Bonhomme et al. (2016) provide somewhat more intuitive conditions that rely, respectively, on sufficient variation in observable covariates that affect choice probabilities and on a sufficient number of conditionally independent measures of the latent state. Both papers establish identification of stationary choice probabilities with panels of length 3. This model provides variation in both dimensions, and so identification is established through either set of results[26].

---

[25]Of course, this size of reduction is not possible, but to simply take it as given provides the most conservative comparison of effect sizes across studies.

[26]The problem can be recast as identification of a finite mixture model with the unobserved state being $(k, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ which has dimension $K \times \mathcal{E}^3$. In this case Kasahara and Shimotsu (2009) require that there be sufficiently many covariates in the state variable S that move choice probabilities across types. Recall that this model has a very high dimension of S,

Table 10

		Number of Arrests	
	(1)	(2)	(3)
Cognitive	−0.194***	−0.038	−0.038
	(0.036)	(0.076)	(0.079)
Behavioral	-0.076**	-0.147***	-0.148***
	(0.033)	(0.054)	(0.053)
Mother's Ed.			0.001
			(0.022)
Constant	0.649***	0.874***	0.859***
	(0.024)	(0.053)	(0.293)
Observations	2,244	877	877
	Implied % Effects of 1 s.d. Improvement		
Behavioral	-11.7	-16.1	-16.1
Note:		*p<0.1; **p<0.05; ***p<0.01	

This table reports estimates of the relationship between total arrests in young adulthood for CDS children and measures of cognitive and behavioral skills. Estimates are calculated by two stage least squares using the first factor score of each skill in the outcome equation and the second factor score as instruments (see section 4.1). Specification (1) shows results for all non-missing observations of CDS children, while Specifications (2) and (3) use only the chosen sample for the paper.

With choice probabilities and transitions in hand, the parameters that govern utility for each type are identified by the parametric restrictions that map them to choice probabilities.[27]

D.4 Estimation of First Stage Parameters

The expectation maximization algorithm provides a tractable approach for maximizing the likelihood. In this algorithm, a current guess of the model parameters produces posterior weights $q(k, \varepsilon | X_m, \hat{\Theta})$[28] and which form a weighted log-likelihood to estimate:

$$
\begin{aligned}
\hat{\Theta}' = \arg\max_{\Theta} \sum_m \sum_k \sum_{t=1}^{T_m} \sum_{\varepsilon_t} q(k, \varepsilon_t | X_m, \hat{\Theta}) \times \Bigg(& \log(\Pi_K(k)) \\
+ \mathbf{1}\{t=1\} \log(\Pi_W^o(\epsilon_1)) + \mathbf{1}\{t < T_m\} & \sum_{\varepsilon_{t+1}} q(k, \varepsilon_{t+1} | k, \varepsilon_t, \hat{\Omega}) \log(\Pi_W(\varepsilon_{t+1} | \varepsilon_1)) \\
+ \sum_{d=1}^{D} \mathbf{1}\{D_{m,t} = d\} & \log(P_d(S_{m,t}; \Theta)) \\
+ \mathbf{1}\{H_{m,t} > 0\} \log \Bigg[\frac{1}{\sigma_W} \phi & \left(\frac{\log(W_{m,t}^o) - \gamma_{W,0,k} - \gamma_{W,1} Age_{m,t} - \varepsilon_t}{\sigma_W} \right) \Bigg] \\
+ \mathbf{1}\{y(m,t) \in \{1997, 2002\}\} \log \Bigg[\frac{1}{\sigma_\tau} \phi & \left(\frac{\log(\tau_{m,t}^o) - \log \left(\frac{\alpha_{\theta,k} \sum_{f \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f,t)}}{\alpha_l + \alpha_{\theta,k} \sum_{f \in \mathcal{A}(m,t)} \Gamma_{\tau,a(m,f,t)}} (112 - H_{m,t}) \right)}{\sigma_\tau} \right) \Bigg] \Bigg)
\end{aligned}
$$
(47)

with iteration performed until convergence. For computational convenience, the value functions $\nu(S_{m,t})$ are also treated as parameters to be estimated, and the dynamic constraints that relate them to utilities are iteratively imposed inside the algorithm. The augmented EM algorithm proceeds as follows, given an initial guess of the parameters Θ^0 and ν^0:

1. E-Step: given Θ^n, calculate choice probabilities and posterior weights over the unobserved states.

2. VFI-Step: calculate ν^n using (8), with Θ^n.

---

including benefit generosity, earnings disregards, the number and age of all children, and tax code parameters, all of which move choice probabilities differently across types, providing identification through this channel. Bonhomme et al. (2016, 2017) require three conditionally independent measures of the latent state, which in this case is provided by setting the state variable as (k, ε_t) and taking observations of earnings (E_{t-1}, E_t, E_{t+1}). Since earnings are continuously distributed, a partioning exists to satisfy rank conditions for each conditional earnings distribution.

[27]A previous version of this paper took a large panel approach to identification, showing sufficient within-panel variation in wages was sufficient to identify mother m's preferences as T_m grew large. This motivated the use of a score criterion with which a clustering approach was taken, following Bonhomme and Manresa (2015). Details are available for this method upon request.

[28]Calculating posterior weights can be made tractable by exploiting the Markov structure of the model using the Baum-Welch algorithm.

3. M-Step: given posterior weights and ν^n, choose Θ^{n+1} to maximize the likelihood in (47), where values ν^n are used to calculate choice probabilities.

4. Check convergence criteria on Θ and ν. Terminate if satisifed, otherwise return to (1).

D.5 Production Estimates

D.5.1 Deriving Moment Conditions

Skills in the PSID-CDS are measured five years apart. Iterating equation (17) forward 5 periods and collecting endogenous residual terms gives the conditional expectation:

$$\mathbb{E}[\log(\theta_{m,f,t+5})|X_m, \theta_{m,f,t}, k] = \sum_{s=0}^{4} \delta_\theta^{4-s} \Big(\delta_{x,a(m,f,t+s)} \mathbb{E}[\log(Y_{m,t+s})|X_m, k]$$
$$+ \delta_{\tau,a(m,f,t+s)} \mathbb{E}[\log(112 - H_{m,t+s})|X_m, k] \Big) + \delta_\theta^5 \log(\theta_{m,f,t}) + \tilde{\mu}_k(a, \boldsymbol{a}). \quad (48)$$

Taking a conditional expectation deals with the fact that some observations of $Y_{m,t}$ and $H_{m,t}$ are missing either due to non-response or missing survey years. The first stage model estimates allow such a calculation. Now collect all of the production parameters into $\boldsymbol{\delta} = (\delta_x, \delta_\tau, \delta_\theta, \tilde{\boldsymbol{\mu}})$ and define the function $\rho_{k,t}$ as:

$$\rho_{k,t}(X_m, \log(\theta_{m,f}); \boldsymbol{\delta}) = \log(\theta_{m,f,t+5}) - \sum_{s=0}^{4} \delta_\theta^{4-s} \Big(\delta_{x,a(m,f,t+s)} \mathbb{E}[\log(Y_{m,t+s})|X_m, k]$$
$$- \delta_{\tau,a(m,f,t+s)} \mathbb{E}[\log(112 - H_{m,t+s})|X_m, k] \Big) - \delta_\theta^5 \log(\theta_{m,f,t}) - \tilde{\mu}_k(a, \boldsymbol{a}) \quad (49)$$

The linearity of ρ implies that we can substitute the factor scores in (14) for unobserved skills:

$$\rho_{k,t}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta}) = \rho_{k,t}(X_m, \log(\theta_{m,f,t}); \boldsymbol{\delta}) + \xi_{\theta,m,f,t+5,1} - \delta_\theta^5 \xi_{\theta,m,f,t,1}$$

and so by the model and measurement assumptions:

$$\mathbb{E}[\rho_{k,t}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta})|X_m, \hat{\theta}_{m,f,2,t}, k] = \mathbf{0}$$

This expression follows the control function identification approach: so long as k is known, a suitably flexible function of type and family composition can control for the endogenous residual in the skill outcome equation. Under this assumption, all observations of income and hours observed in X_m are potentially valid instruments. Notice also that the use of the second factor score $\hat{\theta}_{m,f,t,2}$ is required in order to handle measurement error, as in Cunha and Heckman (2008).

The control function approach is quite strong. It requires that all forms of unobserved heterogeneity that shape skill formation are a function of a set of discrete types which are invertible in the long-run using observed wages and behavior. A weaker requirement is that any unobserved heterogeneity that cannot be inverted from the model is independent of the policy environment variables, Z_m:

$$\mathbb{E}[\rho_{k,t}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta})|Z_m, \hat{\theta}_{m,f,t,2}, k] = \mathbf{0}.$$

Note that the function $\rho_{k,t}$ returns a vector with two entries (one for each skill), $\rho_{k,t} = [\rho_{k,t,C},\ \rho_{k,t,B}]'$.

The EM algorithm from the first stage delivers a posterior distribution over types for each mother, $q(k|X_m, \hat{\Theta})$, which can be used in the moment condition. According to Bonhomme and Manresa (2015), one could simply apply the moment condition using the mostly likely type according to these posteriors, with classification error disappearing asymptotically. Alternatively, the posteriors can be used to integrate out unknown type:

$$\rho_t(X, \theta; \boldsymbol{\delta}, \Theta) = \sum_k \rho_{k,t}(X; \theta, \boldsymbol{\delta}) q(k|X, \Theta)$$

where it is easy to verify that

$$\mathbb{E}[\rho_t(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta}, \Theta)|X_m, \hat{\theta}_{m,f,2}] = \mathbf{0}$$

at the true values of the parameters $(\boldsymbol{\delta}, \Theta)$. The two sets of moment conditions are therefore:

$$\mathbb{E}\left[\begin{array}{c} \rho_{t,C}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta})\frac{\partial\rho_{t,C}(X_m, \hat{\theta}_{m,f,2}; \boldsymbol{\delta}^0)}{\partial\boldsymbol{\delta}} \\ \rho_{t,B}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta})\frac{\partial\rho_{t,B}(X_m, \hat{\theta}_{m,f,2}; \boldsymbol{\delta}^0)}{\partial\boldsymbol{\delta}} \end{array}\right] = \mathbb{E}[g^{\text{model}}_{m,f,t}(\boldsymbol{\delta})] = \mathbf{0} \qquad (50)$$

$$\mathbb{E}\left[\begin{array}{c} \rho_{t,C}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta})\frac{\partial\rho_{t,C}(Z_m, \hat{\theta}_{m,f,2}; \boldsymbol{\delta}^0)}{\partial\boldsymbol{\delta}} \\ \rho_{t,B}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta})\frac{\partial\rho_{t,B}(Z_m, \hat{\theta}_{m,f,2}; \boldsymbol{\delta}^0)}{\partial\boldsymbol{\delta}} \end{array}\right] = \mathbb{E}[g^{\text{iv}}_{m,f,t}(\boldsymbol{\delta})] = \mathbf{0} \qquad (51)$$

where $\boldsymbol{\delta}^0$ is any initial guess of the parameter vector. Subject to rank conditions, any value of $\boldsymbol{\delta}^0$ is a valid choice. In practice, a non-linear least squares routine delivers this initial guess:

$$\boldsymbol{\delta}^0 = \arg\min \sum_m \sum_f \sum_j \rho_{t,j}(X_m, \hat{\theta}_{m,f,1}; \boldsymbol{\delta}, \hat{\Theta})^2.$$

Each estimator $j \in \{\text{model}, \text{iv}\}$ is the solution to:

$$\hat{\boldsymbol{\delta}}^j = \arg\min_{\boldsymbol{\delta}} \left(\sum_{m=1}^M \sum_f \sum_{1997,2002} g^j_{m,f,t(m,1997)}(\boldsymbol{\delta})\right)' \mathbf{W} \left(\sum_{m=1}^M \sum_f \sum_{1997,2002} g^j_{m,f,t(m,1997)}(\boldsymbol{\delta})\right)$$

For each estimator, δ^0 is also used to form a weighting matrix using the inverse of the moment variance:

$$\mathbf{W} = \widehat{\mathbb{V}[g^j_{m,f,t}(\boldsymbol{\delta}^0)]}^{-1}.$$

In practice, when forming the instruments for the moment conditions, some components of the partial derivative must be dropped due to rank conditions. Since the moment condition ρ is linear in the explanatory variables, the control function $\mu_k(a, \boldsymbol{a})$ can first be differenced out by de-meaning all other variables in the outcome equation by age and type. For any variable W this is achieved using the posterior weights $q_{k,m}$ as:

$$\tilde{W}_{m,f} = W_{m,f} - \sum q_{k,m}\overline{W}_{k,a(f)}, \qquad \overline{W}_{k,a} = \frac{\sum_{m,f} \mathbf{1}\{a(f) = a\}q_{k,m}W_{m,f}}{\sum_{m,f} \mathbf{1}\{a(f) = a\}q_{k,m}}.$$

Additional controls for the age of other children in the household ($\boldsymbol{a}$) were found not to matter for estimates. Given estimates, the typical sandwich variance formula delivers standard errors, with a correction for the variance contribution coming from the first stage estimate of $\hat{\Theta}$. In practice, this contribution is small.

D.5.2 Quasi-Bayesian Routine

The quasi-likelihood is computed as:

$$L_M(\boldsymbol{\delta}) = -\frac{M}{2} \left(\sum_{m=1}^{M} \sum_{f} \sum_{1997,2002} g_{m,f,t(m,1997)}^{\text{model}}(\boldsymbol{\delta}) \right)' \mathbf{W}^* \left(\sum_{m=1}^{M} \sum_{f} \sum_{1997,2002} g_{m,f,t(m,1997)}^{\text{model}}(\boldsymbol{\delta}) \right)$$

where $\mathbf{W}^*$ is the inverse of the covariance of the moment conditions evaluated using the GMM estimate $\boldsymbol{\delta}^{\text{model}}$ from the prior section. The priors π are:

$$\log(\delta_{a,i,j}) \sim \mathcal{N}(\log(0.1), 4) \qquad \log(\delta_{\theta,j,j'}) \sim \mathcal{N}(\log(0.5), 4) \qquad i \in \{x, \tau\}, \ j, j' \in \{C, B\}.$$

A Metropolis Hastings algorithm provides a random sample from this posterior distribution via a Markov Chain with 1,000,000 draws using a standard normal proposal density multiplied by a tuning parameter, *temp*. A value of $temp = 0.06$ achieves a mean acceptance probability between 30 and 60%. From this chain, every 100th draw after a burn-in period of 100,000 forms the sample of the posterior distribution for all results in the paper.

D.5.3 Benchmarking Production Parameters

Using the year 2000 as a base year, the *ceterus paribus* effect is calculated for the oldest CDS child in each household after an additional \$1000 in annual income. Letting a_m^* be the age of the oldest CDS child for mother m in the year 2000, this is:

$$\text{Ceterus Paribus Effect} = \frac{1}{M} \sum_{m} \delta_{x,a_m^*} \mathbb{E} \left[\log \left(\frac{Y_{mt} + 1000}{Y_{mt}} \right) \right].$$

To repeat a point made in the introduction, these ceterus paribus effects are only partially informative and do not include any potential effect of the transfer on labor supply, which the model implies will also have some effect on skills. Therefore, calculation of the *endogenous effect* of the \$1000 transfer proceeds by first adding the \$1000 transfer to all household budget sets in the year 2000, and re-calculating choice probabilities given the new budget. Let the superscript $\tilde{\ }$ indicate random variables under this counterfactual, giving:

$$\text{Endogenous Effect} = \frac{1}{M} \sum_{m} \mathbb{E} \left[\delta_{x,a_m^*} \log \left(\frac{\tilde{Y}_{mt}}{Y_{mt}} \right) + \delta_{\tau,a_m^*} \log \left(\frac{112 - \tilde{H}_{mt}}{112 - H_{mt}} \right) \right].$$

D.6 Robustness and Specification Tests

D.6.1 Testing Investment Behavior

While the combination of Cobb-Douglas technology and log preferences produces an important simplification of the model solution, it also places severe restrictions on mother's investment behavior,

which is testable. Recall that time investment for mother m on child k at time t is: [29]

$$\tau_{mkt} = \frac{\alpha_{\theta,m}\Gamma_{a_{k,t}}}{\alpha_{l,m} + \alpha_{\theta,m}\sum_{j=1}^{K_m}\Gamma_{a_{j,t}}} \times (L - H_{m,t})$$

Taking a log-transformation of this specification, and allowing for measurement error in observed time use, we can decompose time investment into an age-specific constant, $\gamma_{a_{k,t}}$, a mother-year specific term μ_{mt}, and measurement error:

$$\log(\tau_{mkt}^o) = \log\left(\frac{\alpha_{\theta,m}}{\alpha_{l,m} + \alpha_{\theta,m}\sum_{j=1}^{K_m}\delta_{\tau,a_{j,t}}\Gamma_{a_{j,t}}} \times (L - H_{m,t})\right) + \log(\delta_{\tau,a_{k,t}}\Gamma_{a_{k,t}}) + \epsilon_{m,k,t}$$

$$= \mu_{m,t} + \gamma_{a_{k,t}} + \epsilon_{m,k,t}. \tag{52}$$

If preferences —given here by the pair $(\alpha_{l,m}, \alpha_{\theta,m})$ —could be identified, then this restriction can be tested directly by including measures of skills in a regression equation based on equation (52). An alternative test that does not rely on exact identification of preference parameters exploits the existence of sibling pairs, which allows the mother-time fixed effect $\mu_{m,t}$ to be differenced out in a fixed-effects approach. Following this approach, the following specification is estimated:

$$\log(\tau_{mkt}^o) = \mu_{m,t} + \gamma_{a_{k,t}} + \beta_1 \text{LW}_{kt} + \beta_2 \text{AP}_{kt} + \epsilon_{m,k,t}$$

using sibling pairs to identify the mother-time fixed effect $\mu_{m,t}$ and age dummies to estimate the term $\gamma_{a_{k,t}}$. Under the null hypothesis that the model has been correctly specified, the estimands β_1 and β_2 are equal to zero. This approach to testing is appealing for two reasons. It has an exactly coherent interpretation within the model (the null hypothesis is structurally motivated), but it also produces the familiar sibling-pair design, which seeks to control for unobserved family level factors (in this case, the preference for child investment interacted with the age-composition of the family) by exploiting within-family variation in skill endowments. Table 11 reports estimates of this specification, using both the preferred measure of active time, as well as an alternative measure of time investment that includes all activities when the mother is around (total time). In order to handle any attenuation bias caused by measurement error in skills, Table 11 also reports results when Letter-Word (LW) scores are instrumented for using Applied Problems (AP), and Externalizing Behavior Scores (BPE) are instrumented for using Internalizing Behavior Scores (BPN). These specifications do not find any evidence of an impact of cognitive or behavioral skills on time investment, and the null hypothesis that $\beta_1 = \beta_2 = 0$ is not rejected. Furthermore, the estimates for these coefficients are fairly robustly and reasonably precisely centered at zero. Given evidence here and elsewhere (Cunha and Heckman, 2008; Cunha et al., 2010) that current skills are complementary with investments, this evidence suggests that the current model fits the data better than alternatives that do not allow for this separability.

---

[29]Appendix C shows that this specification is robust to allowing for public and private time use categories in a Cobb-Douglas specification.

Table 11: Testing for effects of skill on investment behavior

Specification:

$$\log(\tau^o_{mkt}) = \mu_{mt} + \gamma_{a_{kt}} + \beta_1 \text{LW}_{mkt} + \beta_2 \text{BPE}_{mkt} + \epsilon_{mkt}$$

	Active Time		Total Time	
	OLS	IV	OLS	IV
LW	0.002	−0.065	0.017	0.001
	(0.040)	(0.079)	(0.027)	(0.049)
BPE	−0.008	−0.017	−0.002	−0.014
	(0.012)	(0.024)	(0.009)	(0.017)
Age Dummies	✓	✓	✓	✓
Mother × Year FE	✓	✓	✓	✓
Observations	1,463	1,437	1,549	1,522
R^2	0.100	0.086	0.073	0.061

This table displays the results of the specification test of the model using a sibling design. No specification can reject the null hypothesis that $\beta_1 = \beta_2 = 0$. Applied problems (AP) and internalizing behaviors (BPN) scores are used as instruments in IV specifications. Standard errors are clustered at the family level.

D.6.2 Testing Returns to Experience

This section tests for returns to experience by examining whether accumulated labor market experience can predict any of the difference between observed wages and those forecast by the model. To perform this test, first define the residual:

$$v_{m,t} = \log(W_{m,t}^o) - \mathbb{E}[\log(W_{m,t})|X_m, \hat{\Theta}].$$

The model also offers a measure of labor market experience:

$$\text{Exp}_{m,t} = \mathbb{E}\left[\sum_{s=1}^{t} H_{m,s}|X_m, \hat{\Theta}\right]$$

in which missing observations of $H_{m,t}$ are able to be imputed. A testable restriction of the model is that

$$\mathbb{E}[v_{m,t}|\text{Exp}_{m,t}] = 0.$$

Table 12 tests the restriction by analyzing regression coefficients of the residual v under a number of specifications, finding no evidence of a significant and positive return to labor market experience. Furthermore, the power of labor market experience to predict wages beyond the estimated model's forecast appears to be fairly tightly bounded around zero.

D.6.3 Testing for Substitution Patterns

Define the residual $v_{\phi,m,t}$ as the difference between the observed fraction of time at home spent with all CDS children and the prediction formed from the model:

$$v_{\phi,m,t} = \phi_{m,t} - \mathbb{E}[\phi_{m,t}|X_m, \hat{\Theta}].$$

The model implies the restriction that:

$$\mathbb{E}[v_{\phi,m,t}|\log(Y_{m,t}), \log(112 - H_{m,t})] = 0.$$

If the data exhibit patterns of investment consistent with time and money being substitutes or complements, this restriction will be violated. For example, if time and money are complements, one may expect a positive relationship between the model residual v_ϕ and net household income. Table 13 tests for any deviations from the model by regressing v_ϕ on income and hours at home.

This testing approach finds no statistically significant evidence of substitution patterns that deviate from the linear investment rules derived in the model. Point estimates suggest however that increases in time at home may lead to reductions in the fraction of time spent with children, consistent with time and money investment being complements, as found by Caucutt, Lochner, Mullins, and Park (2020). The analysis here suggests that such an effect is not a statistically salient feature of the data for this population and that the model does a good job of capturing the first order effects of work decisions on child outcomes.

Table 12: Test for Returns to Experience

Specification:

$$v_{m,t} = \beta_0 + \beta_1 \text{Exp}_{m,t} + \epsilon_{mt}$$

	(1)	(2)	(3)
Exp	-0.001	0.002	-0.00005
	(0.001)	(0.001)	(0.001)
Individual FE	-	✓	✓
Age FE	-	-	✓
Observations	6,058	6,058	6,058
R^2	0.0001	0.0002	0.015
Note:			*p<0.1; **p<0.05; ***p<0.01

This table displays the results of the test for returns to experience in the data using the model's prediction of log-wages and accumulated experience. No specification can reject the null hypothesis that $\beta_1 = 0$. Standard errors are clustered by mother.

Table 13: Test for Substitution Patterns in Investment

Specification:

$$v_{\phi,m,t} = \beta_0 + \beta_1 \log(Y_{m,t}) + \beta_2 \log(112 - H_{m,t}) + \epsilon_{m,t}$$

	(1)	(2)
$\log(Y_{m,t})$	0.004	0.037
	(0.011)	(0.033)
$\log(112 - H_{m,t})$	-0.137	-0.412
	(0.129)	(0.427)
Constant	0.594	
	(0.645)	
Observations	1,237	1,237
Mother FE	-	✓
R^2	0.007	0.031
Note:	*p<0.1; **p<0.05; ***p<0.01	

This table displays the results of the test for substitution patterns in time investment using the model's prediction of the fraction of time at home spent with children. Neither specification can reject the null hypothesis that $\beta_1 = \beta_2 = 0$. Standard errors are clustered by mother.

E Optimal Policy

For convenience, begin by restating the first order conditions for the problem:

$$\sum_s \pi(s,e)\left\{ P(s,e)\left[\frac{\mu(s,e)\tilde{\alpha}_C(s) + \boldsymbol{\lambda}\tilde{\delta}_x(s)}{\boldsymbol{y}(e)} - \boldsymbol{\lambda} \right]\right.$$

$$\left. + \boldsymbol{\lambda}\frac{\partial P(s,e)}{\partial[u_1(\boldsymbol{y}(e),s) - u_0(\boldsymbol{y}(0),s)]}\frac{\tilde{\alpha}_C(s)}{\boldsymbol{y}(e)}\left[e + \boldsymbol{y}(0) - \boldsymbol{y}(e) + \mathcal{D}(s,e)\right] \right\} = 0 \quad (53)$$

$$\sum_{s,e} \pi(s,e)\left\{ (1 - P(s,e))\left[\frac{\mu(s,e)\tilde{\alpha}_C(s) + \boldsymbol{\lambda}\tilde{\delta}_x(s)}{\boldsymbol{y}(0)} - \boldsymbol{\lambda} \right]\right.$$

$$\left. - \boldsymbol{\lambda}\frac{\partial P(s,e)}{\partial[u_1(\boldsymbol{y}(e),s) - u_0(\boldsymbol{y}(0),s)]}\frac{\tilde{\alpha}_C(s)}{\boldsymbol{y}(0)}\left[e + \boldsymbol{y}(0) - \boldsymbol{y}(e) + \mathcal{D}(s,e)\right] \right\} = 0 \quad (54)$$

Multiplying the first equation through by $\boldsymbol{y}(e)$ and the second through by $\boldsymbol{0}$, then summing the first equation over e, then combining equations results in:

$$\sum_s \pi(s,e)(\mu(s,e)\tilde{\alpha}_C(s) + \boldsymbol{\lambda}\tilde{\delta}_x(s)) = \boldsymbol{\lambda}\sum_{s,e} \pi(s,e)(\boldsymbol{y}(0) + P(s,e)\boldsymbol{y}(e) - \boldsymbol{y}(0))$$

This equation can be rearranged into (27) in the main text. Using the definition of η, (25) can be rearranged into:

$$\sum_s \pi(s,e)P(s,e)\left\{ \left[\frac{\mu(s,e)\tilde{\alpha}_C(s)}{\boldsymbol{\lambda}} + \tilde{\delta}_x(s) - \boldsymbol{y}(e)\right] + \eta(s,e)\left[e + \boldsymbol{y}(0) - \boldsymbol{y}(e) + \mathcal{D}(s,e)\right] \right\} = 0$$

Pulling $\boldsymbol{y}(e)$ to the other side of the equation, applying the definition of $\boldsymbol{w}$ and dividing through by $\sum_s \pi(s,e)P(s,e)$ results in:

$$(1 + \mathbb{E}[\eta(s,e)|d = 1,e])\boldsymbol{y}(e) = \mathbb{E}[\boldsymbol{w}(s,e) + \eta(s,e)(e + \boldsymbol{y}(0) + \mathcal{D}(s,e)|e,d = 1].$$

Notice that when $\eta(s,e) = 0$, this reduces to the first-best allocation:

$$\boldsymbol{y}^*(e) = \mathbb{E}[\boldsymbol{w}(s,e)|e,d = 1].$$

And rearranging (27) gives:

$$\boldsymbol{y}^*(0) = \frac{\mathbb{E}[\boldsymbol{w}(s,e)] - \sum_{s,e} \pi(s,e)P(s,e)\boldsymbol{y}*(e)}{1 - \mathbb{E}[d]} = \mathbb{E}[\boldsymbol{w}(s,e)|d = 0].$$

Substituting in the definition of $\boldsymbol{y}^*(e)$ and rearranging the equation above results in (28) in the main text. To derive equation (29), note that:

$$\boldsymbol{w} = y(0) - \sum_{s,e} \pi(s,e)P(s,e)(\boldsymbol{y}(e) - \boldsymbol{y}(0))$$

and so (28) can be written as:

$$\boldsymbol{y}(e) - \boldsymbol{y}(0) = \frac{\sum_{e'} \pi(s,e')P(s,e')(\boldsymbol{y}(e') - \boldsymbol{y}(0)) + \eta(e + \mathcal{D}(s,e))}{1 + \eta}.$$

Multiplying through by $\pi(s,e)P(s,e)$ and summing over e gives (29) in the main text.

E.0.1 On the Existence of Work Credits

The optimal policy is defined as having a "work credit" if

$$\lim_{e \to 0} \boldsymbol{y}(e) - \boldsymbol{y}(0) > 0.$$

A work credit necessarily implies that $\boldsymbol{y}(e) > e$ if e is sufficiently close to zero earnings and is therefore equivalent to the optimal policy featuring earnings subsidies. Even without incentive constraints $(\eta(s,e) = 0)$, the optimal policy may trivially feature work credits if

$$\lim_{e \to 0} \mathbb{E}[\boldsymbol{w}(s,e)|e, d = 1] > \mathbb{E}[\boldsymbol{w}(s,e)|d = 0]$$

and so let us set $\boldsymbol{w}$ to be constant for the remainder of this analysis. Equation (27) can be rearranged to:

$$\boldsymbol{y}(e) - \boldsymbol{y}(0) = \frac{\boldsymbol{w} - \boldsymbol{y}(0) + \mathbb{E}[\eta(s,e)(e + \mathcal{D}(s,e)|e, d = 1]}{1 + \mathbb{E}[\eta(s,e)|e, d = 1]}.$$

In the standard case without skill formation, the limit at $e = 0$ becomes:

$$\text{work credit} = \frac{\boldsymbol{w} - \boldsymbol{y}(0)}{1 + \lim_{e \to 0} \mathbb{E}[\eta(s,e)|e, d = 1]}.$$

So in this case a work credit is guaranteed in any case in which $y(0)$ is less than average net income among those working. Reintroducing skill formation gives the work credit formula:

$$\text{work credit} = \frac{\boldsymbol{w} - \boldsymbol{y}(0) + \lim_{e \to 0} \mathbb{E}[\eta(s,e)\mathcal{D}(s,e)|e, d = 1]}{1 + \lim_{e \to 0} \mathbb{E}[\eta(s,e)|e, d = 1]}$$

which provides three conjectures. First, work credits appear to be increasing in δ_x at the bottom of the income distribution and decreasing in δ_τ, since these parameters have opposite effects on $\mathcal{D}$. Second, if time is sufficiently important relative to money in skill formation, this would eliminate the presence of a work credit that would otherwise have been optimal in the standard case. In such a case, it is optimal for the planner to offer the worker *less* for working compared to not working. Third, the size of this work credit (or work penalty if δ_τ is sufficiently large) is increasing in the semi-elasticity of participation at the bottom of the income distribution.