Education Savings Accounts, Parent Contributions, and Education Attainment[1]

Michael D. S. Morris[2]

September 2005

Preliminary: please do not quote, cite or circulate without permission.

ABSTRACT

This paper presents a multi-period, dynamic programming model of household choices on savings, consumption, having children and helping to fund children's education. Data from the National Longitudinal Survey young women cohort are used to estimate the parameters of the model. The full structural model is estimated using a simulated maximum likelihood procedure utilizing nonparametric kernel estimators to construct the densities of the likelihood. The estimated model is able to match the general trends in the NLS data, particularly as related to the interaction between children, savings and spending on education. The parameter estimates indicate that the amount that parents choose to contribute to a child's education has a strong impact on the probability that a child attains a college degree, as does the level of education of the parents. Using the estimated model, policy experiments are performed to look at the impact of the creation of tax-free education savings accounts on parental contributions toward education and the education attainment of children. The experiments indicate that accounts do increase contributions and attainment, and the increase is actually greater in lower income households. When compared to other baseline policies such as grants and tax credits, the ESA accounts are more effective than tax credits, but less broad than universal grants.

[1] I would like to thank participants at the 2004 WEAI meetings and 2004 EEA meetings for their comments along with Petra Todd, Jere Berhman and especially Ken Wolpin.
[2] Assistant Professor of Economics, University of New Orleans, Department of Economics and Finance, 2000 Lakeshore Drive, New Orleans, LA 70148; email: mdmorris@uno.edu.

1. Introduction

In the U.S., it is common for parents to help pay for their children's college education and for many families this is a major expense that they face and must plan for. The National Center for Education Statistics (NCES) reports that the average full price (including tuition, room, board and other expenses as reported by universities) for a private 4-year school was $23,600 per year for the 2000-2001 academic year (Berkner, Berker et al. 2002). The cost for a public 4-year school was still $12,600. Furthermore, these costs have been rising rapidly. The National Center for Public Policy and Higher Education estimates that between the 2001-2002 academic year and the 2002-2003 academic year, tuition levels have increased by 7 percent nationally, and by as much as 24 percent in Massachusetts and 20 percent in Texas (Kronholz 2003). Certainly some of this is paid for by student aid through grants, scholarships, and subsidized loans. Still, the NCES reports that 45 percent of students receive no type of aid, either from grants, subsidized loans, or otherwise. McPherson and Shapiro (1991) estimate that for those receiving such aid, the aid only covers 75% of the cost of public schooling and 50% of the cost of private schooling. As such, it does fall upon the individual and their family to help pay for college expenses. Using data gathered by the NCES, Choy and Henke (1992) find that in 1989, 67% of parents contributed to their children's college education and the average annual amount was $3,900 ($5,240 in 1999 dollars, adjusted using the CPI).

Given the high costs of education, the fact that most children need support from their parents to attain a college degree, and the importance of a college degree in determining future earnings and productivity, there has been a push to help parents pay for their children's college education with the introduction of a variety of policies not only aiming to lower the cost of a degree for some children, but also aiming to increase the amount that parents contribute toward the college education of their children. Clearly contributions toward education are contingent on having children in the first place, and if parents do save or adjust their spending behavior over time to help make contributions, then to fully examine the impact of policies on parental contributions it is necessary to consider parents dynamic decisions on fertility, savings and college spending together.

This paper estimates a structural dynamic programming model of household decisions on having children, saving and making transfers to children in the form of educational funding. The model allows for households to face income uncertainty and face some borrowing constraints. In choosing to have children, households gain in utility, but incur costs to raise children. Additional utility is gained from having children receive a college degree, and parents can influence this by offering to help pay for college. The model is estimated using a simulated maximum likelihood procedure and data from the National Longitudinal Survey (NLS) young women cohort. The estimated model is then used to examine a policy experiment aimed to replicate the creation of tax-free education savings accounts (ESA).

The second section of the paper presents a review of relevant previous studies to give a background upon which the current model is based. The third section of the paper describes the augmented life-cycle model that is used and gives the exact specification that is estimated. The section also discusses the issues and numerical techniques associated with solving the model, and concludes with a description of the simulated maximum likelihood procedure used to estimate the model. Section four describes the data, both in terms of how the variables in the model are constructed and descriptive statistics of the sample used. Section five presents the estimated parameter results and discusses the fit of the estimated model. Section six provides a discussion of using the estimated model to examine policy experiments.

2. Background and Previous Literature

The workhorse for economic research on household inter-temporal savings and consumption decisions is the life-cycle model. A very thorough review of the history of these models and their ability to deal with observed microeconomic facts can be found in Browning and Lusardi (1996), including the enhancements of liquidity constraints and precautionary savings.[3] While these models are certainly useful, it has been pointed out (for example see Browning and Lusardi 1996; Keane and Wolpin 2001) that life cycle models really need to account for a whole range of related behavior including not only consumptions and savings decisions but also occupational (including work, leisure and retirement decisions) and fertility decisions. This becomes particularly important when studying the interrelation between parents choice on having children, saving, and contributing toward their children's education.

The need to examine this becomes more important considering the rising cost of post-secondary education, and the wide variety of policies to help students afford a college degree.

[3] For summaries see Hubbard, Skinner et al. 1994; Browning and Lusardi 1996; Coleman 1998

Most government support in the form of the grants and loans is directly targeted at the student attending school, and not parents, and several studies have tried to measure the extent these programs help increase college attendance. Using data from the introduction of the Georgia HOPE scholarship program as a quasi-experiment, Dynarski (2000; 2002) finds the availability of an additional $1,000 subsidy increases attendance by 4%, a figure she finds consistent with previous findings. Ichimura and Taber (2002) estimate a 4.5% increase in attendance from the availability of a $1,000 subsidy using estimate this using a reduced form estimation derived from the model used by Keane and Wolpin (2001) and use the same NLSY data as the Keane and Wolpin study. As discussed by Keane (2002), Keane and Wolpin actually estimate that a $100 tuition increase (a negative subsidy) would lower enrollment rates of 18-24-year-olds by 1.2%. Interestingly, while finding that borrowing constraints are indeed tight for students (they can not even support one year from borrowing alone), Keane and Wolpin do not find that allowing for easier loan access will increase attendance. Instead, they find that the major impact of reducing the borrowing constraint is in a reduction of working by students. Perhaps more importantly, they do find that parental transfers contingent on college attendance are not only prevalent, but do significantly increase the educational attainment of children.

Some recent programs have been created that target parents directly in trying to help them pay for the college education of their children, and potentially increase the amount they spend on education. Toward this end, the Taxpayer Relief Act of 1997 introduced two tax credits and an educational savings account. The Hope Scholarship Credit is a tax credit for 100% of the first $1,000 of qualified tuition expenses and 50% of the next $1,000 of qualified tuition expenses[4], and the credit can be claimed for expenses for yourself or a dependant. The Lifetime Learning Credit is similar except that it is available for any year of education (not just the first two years), but is only for 20% of up to $10,000 in expenses[5].

The act also set up an educational savings account, now called a Coverdell Account, which is a custodial account for children under 18 for which the earnings and distributions are

[4] This is only available in the first two years of post-secondary education and is phased out at higher incomes ($40,000-$50,000 for individual taxpayers, twice that for married households).

[5] Prior to 2003 the limit was only up to $5,000. Also, the Hope Scholarship Credit can be claimed for each child supported separately while the Lifetime Learning single credit is a per family credit.

tax-free as long as they are used for qualified education expenses[6]. In addition to this, 529 savings plans are available[7]. These are administered by states, so the details can vary, but in general there are two types of 529 plans. The first type is a prepaid tuition plan where money is contributed in a child's name and locks in an associated percentage of college expenses at current tuition rates at one of the state's public universities[8]. The other type of plan is another college savings account where contributions that can grow tax-free, and generally withdrawals are also tax free as long as the money is spent on education[9]. More recently, President Bush has repeatedly discussed the creation new education savings accounts that would replace the Coverdell accounts and allow for $7,500 in contributions per year.

For most of these savings plans, though, there is very little evidence as to how much they will impact the amount of money families spend on education, or the impact it has on college attendance, and the related literature on targeted savings, namely IRA's, gives rather mixed results. The model and work presented here addresses these issues. The choice to have children, the number of children to have and the timing of when to have children are all made as endogenous decisions within the life-cycle framework. Furthermore, families are allowed to choose to make transfers to their children in the form of educational subsidies. Together, this allows for a rich relationship between the number of children to have, how much to provide these children with a form of inter-vivos transfers, and savings and consumption within the life-cycle.

3. **The Model**

3.1 The General Model

The model is a based on a household general dynamic programming problem. Households have a lifespan of *T* periods and face earnings uncertainty throughout their life. In each period households choose how much to consume and save. Younger households, in periods 1 through

---

[6] Originally no more than $500 may was allowed be put into the account per-year and even this is phased out at higher income levels. This amount increased to $2,000 in 2002. Also, the distributions were not originally tax-free, though that also was adjusted in 2002.

[7] Also created in 1997, they are named for the IRS tax code that allows for favorable tax treatment to qualified state tuition programs.

[8] If you choose to use the money somewhere other than one of the state's public schools, you will get the amount you contributed to the account, but often without any earnings.

[9] Specific limits vary by states, but exceed the very low limit of the Coverdell account, though at the federal level all contributions are subject to gift tax laws, which are not triggered until gifts exceed $10,000.

T^*, also choose also choose how many additional children to have in each of those periods. Children are assumed to cost Ψ per child, per period until they move out after M periods.

To make these decisions, households maximize their expected discounted utility. Each household's contemporaneous utility depends on both the level of consumption, the number of children in the family, and additional utility as their children become more educated. A household's utility for any given period, t, is then in general given by

$$U(c_t, n_t, e_{1t}, e_{2t}, ..., e_{Et}, \varepsilon_{ct}, \varepsilon_{nt}, \varepsilon_{et}) \tag{1}$$

where $U(.)$ is the utility function, c_t is consumption at time t, n_t is the total number of children at time t, e_{jt} is the number of children with education level j at time t where j is one of E education levels (j=1,...,E), and ε_{ct}, ε_{nt}, and ε_{et} are taste shocks to consumption, number of children and education levels respectively.

When a child moves out at age $M+1$, the household has the option to contribute money toward his or her college education. At this time, the household chooses a one-time offer, o_t, of a per-year amount to contribute toward the child's college education. Given this offer, the child then attains a certain an education level. From the household's view, a child's education level is a realization of a stochastic process that depends on the amount of support offered by the household, along with other possible observable demographic variables. For the household, letting d_{jt} be the education level child j receives in period t,

$$d_{jt} \sim D(d_{jt} \mid w_{jt}) \tag{2}$$

so that d_{jt} is a realization of the conditional distribution $D(d_{jt}|w_{jt})$ where w_{jt} is a vector including the amount offered by the household to child j, along with other possible factors including a family type effect. The amount the household actually pays toward a child's college education is then a product of the per-year offer made and the number of years of schooling the child attends.

Throughout their lives, households face an uncertain income stream, which may depend on certain household characteristic such as level of education. In any given period t, a household earns a specific amount of income, I_t. The process differs before and after retirement and all households are assumed to retire at age T^{**}. As such, I_t is a realization,

$$I_t \sim \begin{cases} H(I_t \mid Z_t) & \text{if } t < T^{**} \\ H_{ret}(I_t \mid Z_t, I_1, I_2, ..., I_{T^{**}-1}) & \text{if } t \geq T^{**} \end{cases} \tag{3}$$

where Z_t is a vector of household characteristics and H(.|.) and H_{ret}(.|.) are the appropriate distributions of income before and after retirement.

The households' problem, then, is to solve

$$\max_{\{c_1 \ldots c_T, n_1, \ldots, n_T, o_{M+1}, \ldots, o_T\}} E\left[\sum_{t=1}^{T} \delta^{\,t-1} U(c_t, n_t, e_{1t}, e_{2t}, \ldots, e_{Et}, \varepsilon_{ct}, \varepsilon_{nt}, \varepsilon_{et})\right] \tag{4}$$

where δ is the discount rate. The maximization is made subject to

$$c_t = k_t(1+r) + I_t - n_{ht}\Psi - \sum_{n=1}^{N_t} a_{nt} - k_{t+1} \tag{5}$$

where k_t is the level of household assets at time t, r is the real interest rate, n_{ht} is the number of children living at home at time t, a_{nt} is the amount the household pays toward child n's college expenses at time t and the remaining variables are defined previously. The expectation is over the three taste shocks and the realizations of future income and children's educational attainment.

3.2 Model Specification

To implement this model, one needs to detail specifications for the utility function, the random taste shocks, the income and children's educational attainment processes, and to set values for the number of periods, T, the retirement age T^{**}, the age through which households may have children, T^*, the number of periods children live at home, M, and the different levels of education that children may attain. The specification used in this paper will allow the model to be solved and subsequently estimated using the NLS data described in section four.

A period in the model is assumed to last 6 years. Since assets evolve slowly over time the assumption of such a long period should not significantly impact the results[10]. The maximum age, T, is set to 12 giving households an adult lifespan of 72 years (corresponding to ages 18 to 90). Households are assumed to retire at age 66, so T^{**} is set to 9.

Households may choose to have children during any of the first three periods, after which they no longer have children, so T^* is set to 3. This implies that all households have all their children by an age of 36. In addition, households are restricted to having no more than 4 children

[10] As will be seen later in the paper, the 6-year period also closely corresponds with the frequency that asset information is collected in the data and the general biennial data collection for the NLS. Furthermore, the longer periods will greatly reduce the computational complexity of the model solution.

in any one period, and no more than 5 children in total[11]. Children are assumed to move out after 3 periods at home, so M is set to 3. This corresponds to children moving out at age 18. Actually, this assumption just means that the household incurs a cost of raising a child for only 3 periods. In the period a child moves out (the 4th period of a child's life, or ages 19-24), the household has the option to help pay for that child's college education. It is assumed that parents only make these contributions in the period that their children move out[12].

When making an offer of financial support for college education, parents are restricted to make the same offer to all children moving out in a given period. This assumption allows for the identification of the impact that offers make on how much schooling children receive. This is necessary because the data actually only report the amounts parents actually contribute toward their children's education. As such, there is no data on this for children who did not go to college, and yet it is probable that some these children would have received some support from their parents. The assumption here is that those children receive the same offer of support as their siblings attending school within the same period. When looking at the NLS data, this does not seem overly restrictive as within a period, most families do not greatly vary the amount they contribute toward their different children. In fact, 50.5% made contributions that differed by less than $1,000 and almost 20% of contributions differed by less than $50.

The restriction that children born within the same period receive the same offer, while providing identification for the impact of the offer, does imply that if offers positively impact the educational attainment of children, then one should observe on average, when comparing families with the same number of children, that families with more children in college will be making larger contributions toward their each of their children's education. It is not obvious that this should be the case because, for example, a family with two children and both of them in college at the same time may have less money available to support either individually than if just one was attending college. Looking at the NLS data, though, for families with two children born

[11] These restrictions impact only a small portion of the data and significantly ease the computational burden of solving the model. Within the NLS data (the data are described in detail in a later section), only 7.18% of households have children when older than 36, so this is not too stringent a limitation. Furthermore, only 1.94% had more than 4 children in any single six-year period and only 4.27% had more than 5 children. In all, only 10.78% of the NLS sample violates one or more of these three restrictions (i.e. had more than 4 children in a given 6-year period, had more than 5 children in total and/or had children after the age of 36).

[12] This assumption again significantly eases the computability of the solution and does not impact very many households. In fact, in the NLS data only 14.45% of children attending school received support after age 24, and for most of these that support came within the next two years.

within six years of each other, the average per-child, per-year contribution in 1999 dollars is \$5,721 for those with both children attending college, and only \$3,666 for those with just one child in college, which is consistent with the assumed restriction.

The stochastic process that determines the level of education a child receives, that is $d_{tj} \sim D(d_{tj}|w_{tj})$, is assumed to be an ordered probit. The vector of factors, w_{tj}, includes a quadratic in the family contribution offer and the level of the parents' education. It is assumed that the outcome is one of three discrete outcomes $d_{tj} \in \{\text{low, medium, high}\}$. These low, medium and high outcomes correspond to high school or less, some post-secondary education or a 2-year degree, and a bachelor's or higher degree.

The household contemporaneous utility function for any period is assumed to take the form

$$U(c,n,q_m,q_h,\varepsilon_c,\varepsilon_n,\varepsilon_q) = \frac{(c\varepsilon_c)^{1-\gamma}}{1-\gamma} + (\lambda_1 + \lambda_2 n + \varepsilon_n)(n + \lambda_3 c) + (\alpha_1 + \alpha_2 q_m + \varepsilon_q)q_m + (\theta_1 + \theta_2 q_h + \varepsilon_q)q_h \tag{6}$$

which combines a constant relative risk aversion (CRRA) utility function for consumption with additively separate quadratic utility in the number of children, an interaction with the number of children and consumption, and a quadratic in the number of children with some college education and the number of children with a bachelor's or higher degree. The utility parameters for children and education are allowed some heterogeneity with respect to different education levels, and the child parameters are allowed to further vary in the first three periods. The taste shocks are assumed to be iid and distributed

$$\ln(\varepsilon_{ct}) \sim N(0,\sigma_c^2),\ \varepsilon_{\text{nt}} \sim N(0,\sigma_n^2),\ \varepsilon_{\text{qt}} \sim N(0,\sigma_q^2). \tag{7}$$

Household income is assumed to be determined by,

$$I_t = \begin{cases} \exp\{Z_t\beta + \varepsilon_{I\,t}\} & \text{if } t < T^{**} \\ b\dfrac{1}{T^{**}-1}\displaystyle\sum_{t=1}^{T^{**}-1}\int \exp\{Z_t\beta + \varepsilon_{I\,t}\}d\Phi(\varepsilon_{I\,t}) + \varepsilon_{rt} & \text{if } t \geq T^{**} \end{cases} \tag{8}$$

where the characteristic vector Z_t contains dummies for the levels of education or parents interacted with a quadratic in age. This gives a standard semi-log specification for pre-retirement income and post-retirement income is then a certain percentage, b, of the average expected pre-retirement income. The pre-retirement and post-retirement income shocks are assumed to be distributed iid[13],

$$\varepsilon_{It} \sim N(0, \sigma_I^2) \text{ and } \varepsilon_{rt} \sim N(0, \sigma_r^2). \tag{9}$$

While earnings certainly may be argued to show some short-term persistence, the iid assumption is not unreasonable for the longer six-year periods of this model.

Because the impact of marital transitions is not a focus of this paper, marital status is assumed to be constant over the agent's life. Without this assumption, changes from single to married and vice versa will be accompanied by shocks to not only income but to assets as well. In addition, the separate contributions toward their children's education from divorced couples would also need to be modeled. Since the sample or stable single parents is so small, the model here will actually only consider stable, married families. This assumption allows for a simpler model that focuses on the household savings and educational transfer decisions, though it clearly will limit to some extent the inferences drawn from the results. The level of education of the household for earnings consideration is also a constant in this model.

Finally, households are also assumed to face the following constraints:

$$\begin{aligned} &c_t \geq 0, \forall t \\ &k_t \geq 0 \;\; \forall t \\ &n_t + 4 \geq n_{t+1} \geq n_t \;\; \forall t \\ &n_t \leq 5 \;\; \forall t \\ &4n_{4t} o_t \in \left[0, (1 + r_k)k_t + I_t - n_{ht}\psi\right] \;\; \forall t \end{aligned} \tag{10}$$

where n_{4t} is the number of children of age 4 in period t and o_t is the per-year, per-child offer. The first two conditions imply that households have positive consumption and that they cannot take out uncollateralized loans. The third condition restricts households to have no more than four additional children in a period and that the number of children cannot decrease. The fourth restriction is the maximum family size of five children. Lastly, the offer is restricted so that it must be non-negative and, because the offer is binding, the parents are not allowed to offer more than can be covered without taking out an uncollateralized loan.

The decision timing of the model is as follows. At the beginning of each period the household receives a realization of ε_{ct}, ε_{nt}, ε_{qt}, and ε_{It} or ε_{rt}. After the realization of these shocks,

---

[13] While earnings certainly may be argued to show some short-term persistence, the iid assumption is not unreasonable for the longer six-year periods of this model.

households receive income and pay expenses for children living at home. Households then make decisions for the period. In the first three periods, the choices of an amount to save and the number of children to have in that period are made simultaneously. In periods 4-6, if parents have children of college age, the educational support offer and the savings decisions are made sequentially. Parents first choose a per-child, per-year amount to offer their children in educational support. Parents then realize the outcomes of their children's education choices and make the appropriate payments. Finally, the household chooses a level of savings and consumption.

3.3 Solving the Model

The household optimization problem can be rewritten recursively. Let $s \in S$ where S is the state space and s is a specific point within the state space. The state space for the for a household is $S = K \times N \times N \times N \times Q \times Q$ where k is the level of assets, $k \in K \subset R_+$, n_t is the number of children born in period, $n_t \in N = \{0,1,\ldots,4\}$ and q_{jt} is the number of children of education level j in period t, $q_{jt} \in Q = \{0,1,\ldots,5\}$. Since the household education level and marital status do not change, they can be suppressed from the state space and the model solved separately for type. Age is a state variable, but the period subscripts capture that so it is not listed in the vector s.

In every period the household chooses some level of consumption and savings, given their state variables. At the same time, if the household is of the appropriate age, they also choose the number of additional children to add to the family. I will refer to these decisions as the primary problem. In periods 4, 5 and 6, the household may also solve another problem, the offer to children. Sequentially, a household solves the offer problem at the beginning of the applicable period and then solves the primary problem for that period after observing the education outcomes of their children. In all other periods, the household just solves the primary problem. Let $O_t(s)$ be the value of entering the offer problem at age t with state s and let $V_t(s)$ be the value of entering the primary problem at age t with state s. $O_t(s)$ is the solution to

$$\begin{aligned} O_t(s) &= \max_o E_d\left[V_t(s') \mid o, s\right] \quad \text{if have children to make offer to} \\ O_t(s) &= V_t(s) \quad \text{otherwise} \end{aligned} \tag{11}$$

subject to (5)-(10), where the expectation is taken with respect to the educational outcome of the children. $V_t(s)$ can now be written as the solution to

$$\begin{aligned} V_t(s) &= \max_{c,k',n'} \{U(c_t, n_t, q_{mt}, q_{ht}, \varepsilon_{ct}, \varepsilon_{nt}, \varepsilon_{ht}) + \delta E[O_{t+1}(s') \mid s]\} \text{ if t = 3,4,5} \\ V_t(s) &= \max_{c,k',n'} \{U(c_t, n_t, q_{mt}, q_{ht}, \varepsilon_{ct}, \varepsilon_{nt}, \varepsilon_{ht}) + \delta E[V_{t+1}(s') \mid s]\} \text{ otherwise} \end{aligned} \tag{12}$$

subject to (5)-(10) where a " ′ " on a variable indicates the value in the next period. The expectations here are taken with respect to ε_c, ε_n, ε_q and ε_I or ε_r. It is assumed that $V_{T+1}(s) = 0$.

The model can be solved as a finite horizon dynamic programming problem. Given a set of parameter values, the direct way to solve a finite-horizon problem is to solve the problem for every possible combination of state points backwards from the final period to the initial period. However, this method is not a practical option when the number of state points becomes too many or when there are continuous variables in the state space[14]. This is compounded by the fact that multiple integration over the random shocks is necessary at each step to get the expected value functions (EMAX - i.e. the expectations in (11) and (12).)

The specification assumptions used here that limiting the number of children and the number of periods that families can have children help limit the number of combinations of children born at different periods and their educational attainment in the state space to a manageable number. However, since assets are a continuous variable, the solution still cannot be calculated for every state variable combination. Even if assets were to be discretized in a reasonable way, the resulting state space would still be huge.

As a method for dealing with this, Keane and Wolpin (1994) propose using Monte Carlo integration to evaluate the expectations in the EMAX and using least squares to interpolate over a subset of state points for which the model is solved. They offer extensive Monte Carlo simulations showing the effectiveness of this method. The solution technique used here follows the Keane-Wolpin methodology. The model will be solved for every possible combination of the discrete state variables (the number of children in each period and their education levels) but for each of these combinations I will only directly solve for a subset of asset values. The expectations are approximated using Monte Carlo integration and the EMAX functions will be interpolated at the remaining asset levels using least squares. For (12) the interpolation is the

[14] The reason is that there is a "curse of dimensionality" due to the fact that the number of combinations of state points grows too large or even infinite so that computing a solution for each one is not possible.

regression of directly solved EMAX values on the contemporary utility evaluated at the means of the stochastic components and for (11) the optimal offers are regressed on assets (Keane and Wolpin, 1994).

3.4 Estimation

The model is estimated using Simulated Maximum Likelihood[15]. The solution to the dynamic programming problem provides the input into estimating the likelihood. For each individual, this problem is deterministic, but from the economists view it is probabilistic because we do not observe the contemporaneous shocks. As such, a likelihood function can be constructed and estimated based on the outcomes predicted by the model solution at different parameter values.

Several previous studies have utilized a similar estimation procedure for discrete choice models.[16] Since the model in this paper contains continuous choices, the exact same procedure cannot be followed.[17] However, the same concept applies, but instead of probability simulators I use a nonparametric kernel density estimator with simulated samples to construct the elements of the likelihood.

For a single household, *i*, the data provide and observable sequence of state points (and hence the decisions they made). Since the decisions made in the model depend only on the current state variable and exogenous, independently distributed shocks, *i*'s contribution to the likelihood, L_i can be rewritten as a sequence of conditional densities:

$$L_i = \prod_{t=0}^{T-1} f(s_{t+1}^i \mid s_t^i) \tag{13}$$

where *f(.)* is the pdf of s_{t+1} conditional on s_t. The sample likelihood is then calculated as the product of these individual likelihoods. Calculating the likelihood is still a problem because the functional form of *f(.)* is unknown. Given a set of the parameters, however, the model can be solved and therefore a sample of values for s_{t+1} given a value of s_t can be simulated. From this sample, then, a density estimator can be calculated and used to estimate the value of $f(s_{t+1}|s_t)$.

[15] For a good summary of simulation based estimation techniques, see Gourieroux and Monfort (1996).

[16] Simulated probability estimators utilizing simulated samples generated from the model are used to estimate these probabilities and therefore generate the likelihood which is maximized (for examples of such simulators see McFadden 1989; Geweke, Keane et al. 1994; Stern 1994; Keane and Wolpin 1997).

[17] Keane and Wolpin (2001) develop an expanded approach for a model that maintains a discrete choice set but allows for additional continuous and unobserved outcome variables, such as income. They accomplish this with an innovative use of assuming there is measurement error for the unobserved variables.

There is a wide literature on density estimation and techniques for continuous and mixed variable distributions (for a summary see Silverman 1986). Because of the maximization problem involved, a smooth density estimator is important here, and I will use a smooth kernel density estimator. Unfortunately, traditional kernel estimators are notoriously inaccurate and difficult (if not impossible) to implement when used in higher dimension space. However, in the specific model here, it turns out that no more than two observed state space values change between any two observed periods in which the conditional density must be estimated. The remaining values are fixed. As such, it is never necessary to estimate a joint density for more than two variables, which can be done to reasonable accuracy. For example, after period six, the number of children and their education levels are all fixed so just assets change, so just a one-dimensional conditional density of next period assets given current assets is needed. In the education periods, 4-6, the number of children is fixed, but there is an offer decision, and education levels and assets are changing. However, this process is sequential in the model: first an offer is made, then educational outcomes are realized conditional on this offer, and then savings decisions are made. As such, there are really three independent conditional densities for these periods, one of offer given the initial state, the next of educational outcomes given the offer and the initial state, and finally of savings given the initial state and educational outcomes. But again, only a one-dimensional density estimator is needed for the offer and asset updates, and a two-dimensional estimator for the educational outcomes for changes in the number of children with some college or a four-year degree. In the first three periods, the education levels are constant, as are the number of children born in any period except the current one, so only the number of children in the current period and savings are changing. As such, again only a bivariate density estimator is needed for these periods.

When a univariate continuous density estimator is needed, I use a standard Gaussian kernel,

$$K(s, S_i, h) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{s - S_i}{h}\right)^2 \right) \tag{14}$$

and for the smoothing parameter I use Silverman's (1986) plug in value for an optimal window width for the Gaussian kernel,

$$h = 1.06\hat{\sigma}n^{-1/5} \tag{15}$$

where $\hat{\sigma}$ is the standard deviation of the variable in the sample data[18]. For the bivariate densities I use a bivariate product kernel (Scott 1992) which in general takes the form

$$\hat{f}(s_1, s_2) = \frac{1}{nh_1h_2}\sum_{i=1}^{n} K_1(s_1, S_{i1}, h_1)K_2(s_2, S_{i2}, h_2). \tag{16}$$

When the bivariate estimates are needed, they are either for either two discrete variables (educational outcomes when appropriate in periods 4-6) or mixed with a continuous variable (assets) and a discrete variable (new children) in periods 1-3. For the continuous portion I again use a standard Gaussian kernel as described in (14)-(15). For the discrete values, the variables here are not just categorical, but are ordered and the values matter. To take advantage of this, I utilize a variant of the Habbema kernel which was found to be highly effective by Titterington and Bowman (1985) with

$$K(s, S_i, h) = \lambda^{|s-S_i|} \tag{17}$$

where λ is a smoothness parameter. In this kernel smoothness weight, *h*, is set to

$$h = \sum_{j=0}^{J-1} \lambda^j \tag{18}$$

where *J* is the number of discrete distances possible between *s* and other data points in the sample and acts just as an appropriate weight. The amount of smoothing is then controlled by *λ*, which I set to 0.3[19].

The estimation procedure works as follows. First, an initial guess of the parameters is made. The model is solved for these parameters. The value of the likelihood function is constructed using the simulated density estimation just described. The likelihood is checked to see if it is maximized, if it is not, the guess is updated and the procedure repeats itself. The parameters are updated for the maximization using a version of a simplex algorithm originally proposed by Nelder and Mead (1965) and further discussed in Acton (1990).

---

[18] This choice of a value for the smoothing parameter (and slight variants) is also proposed and discussed in Härdle (1991) and Scott (1992). In a comparative simulation study of kernel methods, Bowman (1985) finds that such a plug-in parameter selection, while simple, performed very well even when compared to more advanced smoothing parameter selection methods.

[19] Setting λ=0 results in a histogram and setting λ=1 results in an equal density estimate of $1/J$ for all values. 0.3 is utilized by Titterington and Bowman (1985) and as with their studies, I find that changing this within a reasonable range does not qualitatively impact the estimates. For more discussion on kernel methods for discrete distributions refer to Aitken (1983).

4. The Data

The data are taken from the young women cohort of the National Longitudinal Survey (NLS), which consists of 5,159 women who were 14-24 years old in 1968. Surveys were administered every year from 1968 through 1973 and then basically every other year since then. The sample used here covers the years 1968-1999. While the NLS collects data on a wide variety of topics, the relevant data used in this paper are on assets, income, children, education, marital status, children's education and spending on children's education.

4.1 Sample and variable definitions.

When matching the 6-year period in the model with the data, the first period is matched to ages 18-23 in a household's life, the second ages 24-29 and so on. The age of the household is measured as the age of the women followed in the NLS data.

Comprehensive questions on assets were only asked in 1968, 1971-1973, 1978, 1983, 1988, 1993, 1995, 1997 and 1999. For this paper, I define assets as total net household assets, excluding vehicles.[20] If more than one observation of assets is available within the appropriate age range for a period for a household, the earlier dated value is used because the model treats the assets of a period as the assets available to the household when entering that period.

Income information is collected with every survey, but the level of detail varies some between years.[21] Income in the model is measured as the total income of a woman and her spouse. To fit with the six-year periods, the income amount is calculated as six times the average of the annual income observations within the appropriate age range.

Educational information about the respondents is updated in every survey. Information on the level of education of spouses is also collected. The level of education enters the model in the wage equation and as a determinant of the level of education attained by children in the household. Since the model assumes that the level of education is constant for a household's lifetime, the household education level is measured as the highest level of education attained by either parent in a household. The measured education level of a household is then grouped into

[20] Asset information obtained includes information about housing, mortgages, savings, stocks, bonds, vehicles, business, farm or real estate assets and outstanding debts.

[21] Questions generally cover wages of the respondent and spouse, business or farm income, rental income, unemployment compensation, disability income, welfare, income of other family members and other income. In

one of two categories, those with a college degree, and those without, giving two education "types" of households. The limit to two education categories is because the sample sizes within education groups become rather small when more categories are used.

Marital status is asked in every interview and in most surveys the marital history since the last interview is updated. As with the level of education, marital status is assumed to be constant in the model. I classify a household as married in the model if the woman is married before age 36 and in no subsequent interview after getting married are they no longer married. In similar fashion, someone never married is considered single. However, as previously mentioned, the number of women who were never married and also satisfy the other data assumptions and restrictions is extremely small. Because of this, I will only consider stable, married households.

In every interview, information, including age but not education, is collected for each of the respondent's children living in the household. Furthermore, in 1999 a full child roster was collected which includes children's birth-dates, and level of education. From this information, children can be grouped into being born in the first, second or third period of the model and the appropriate educational attainment category can also be assigned for each child. As previously mentioned, the model assumes households can only have children in the first three six-year periods, and that they have no more than five children, and no more than four within a given period. Only 10.78% of the NLS sample violates one or more of these three restrictions and will not be used. This restriction, while perhaps not a large loss of data, should be kept in mind when considering the types of households for which the results are estimated.

Beginning in 1991 and subsequently in 1993, 1995, 1997 and 1999, the survey includes questions about the college enrollment of children and the amount of financial support parents provide toward college for each child within the past twelve months. The offer is measured as the average annual contribution parents made toward post-secondary education for children born in the appropriate period. Unfortunately, the only data collected on parent's contributions is the biennial question about how much was contributed within the last 12 months. There is not data on the total amount spent on a child's post-secondary education. Some measure of this is needed to update the household assets level for the total amount parents spend to assist their children's

---

some years, separate questions on interest and dividend income, food stamps, alimony and child support, social security and pension income and assistance from relatives were also asked.

education. This amount will be computed as 4 times the offer for those children within the four-year college degree education category, and 2 times the offer for those children within the some college education category. This is not an amount matched to any particular figure in the data, but is used to update the household's evolving assets.

The original NLS young women cohort sample began with 5,159 women of ages 14-24 years old when the initial data collection began in 1968. For analysis with the model in this paper, this set is restricted to those women who remained in the data set through 1999 and reported information on all the necessary and relevant variables. One impact of this is that it somewhat overly excludes the older women in the sample. This is because data on contributions toward their children's education was only gathered starting in 1991. Therefore, some of these women had children already educated before any information on their contributions was collected. However, there is no reason to suspect that these women are significantly different from the younger women in the cohort so this should not be a problem. The sample is further limited to those women who meet the stable marriage criteria outlined above. Lastly, given the size of the sample and the likely measurement error that accounts for extremely high and low reported assets levels, outlier observations were deleted.[22] This leaves a sample of 556 households with 3,139 household-period observations.

4.2 Descriptive Statistics

Table 1 shows some descriptive statistics of the sample, along with those of a prediction from the model which will be discussed later. The "actual" column gives the sample descriptives. When looking at the number of children, less educated households have more children, averaging 2.395 per household versus 2.208 for households with a college degree. Not surprisingly, the timing of when to have children is also very different. College educated households have very few children, only 0.476 on average, during ages 18 through 23. Conversely, this is when households without a college degree have the most children, averaging 1.225 children per household born during ages 18 through 23. At older age ranges, however, it is the college-educated households who have more children.

---

[22] The upper and lower truncation points differed for each period. In total, 21 observations from below and 28 from above were cut.

Table 1 also presents information about the amount of money families spend on their children's college education. For the entire sample, 81.7 percent of households with children attending college helped pay for college. Furthermore, the average per-year contribution made by those who did help financially support their children was $5,230 in 1999 dollars. The percentage or parents who help support their children in college rises to just under 95 percent for households with a college degree. Even for households with no more than a high school degree, the percentage remains high at 67.7 percent. The average amount contributed for those families who helped pay for college is $6,690 for households with a college degree and $3,033 for those without. This is a per-year, per-child figure. Clearly, this indicates a significant amount of money, especially when multiplied over several years and several children.

Looking at the distribution of education attainment for children, overall 21.9 percent receive a four-year degree or more, 35.73 percent attend some college but do not earn a four-year degree and 42.37 percent do not attend college at all. For households with a college degree, the percentage of children earning a four-year degree or more rises to 37.34 percent and only 23.10 percent do not attend college at all. This relationship between parents and children's education has been consistently documented (see Haveman and Wolfe 1995 for a review).

Table 2 presents some more detailed descriptive statistics on the asset accumulation of households at different ages and by differing levels of education. Since the oldest in the cohort are just now reaching their mid-to-late fifties, the accumulation pattern is just increasing through the years. As usual, the asset distribution is largely skewed toward high values as is shown by the much lower median asset values. Ten percent of the sample lists negative asset levels at ages 18-23 but this progressively declines to only 1.89 percent by ages 48-53. Perhaps more significantly, the average value of assets for those with a negative asset level is never very far from zero, ranging from just less than $-1,000 to just under $-3,000. So, while some households do have negative net assets, the percentage becomes more negligible as households age, and the negative amount does not tend to fall much below zero. This is consistent with previous research on the liquidity constraints on households (see for example Keane and Wolpin 2001) and is not too far from the constraint against non-collateralized loans in the model estimated here, particularly after period one.

Figure 1 breaks down household asset accumulation by the number of children parents have over their lifetime. As shown in the top panel, the most striking feature is that households without children households with no children have the lowest assets in early years, but by their late forties, these households have the highest average asset level. The general pattern for households with children is to accumulate assets at a quickening pace until their early thirties, to slow asset accumulation during their thirties, but then resume at a more rapid pace in their mid-forties. Households with 1-3 children have significantly more assets than those with four or five children. This pattern is similarly repeated when separating households with a college degree from those without as is done in the lower two panels of figure 1. The observed asset accumulation pattern is certainly consistent with households saving more to cover the cost of children and then slowing their asset accumulation while incurring expenses associated with children.

5. Results

5.1 Parameter Estimates

The estimated parameters of the model are reported in table 3. There are a total of 48 parameters estimated here, and Figure 2 explicitly shows the parameters in the specification used here. While that is a large number of parameters, the data consist of over 3,139 household-period observations.

The estimated parameters for the children's educational outcome suggest that both parental offers for financing and the parents' education level have a sizeable impact on the outcome. The marginal effects indicated by the ordered probit estimates are nonlinear and themselves functions of the values of parental contributions and education levels.[23] Table 4 shows some of the predicted probabilities for a child's education for different levels of parental contribution and education, along with the marginal impact of an additional dollar on the probabilities. For evaluating the impact of the level of parents' education, the marginal impact can be seen by comparing the change in the predicted probabilities. Evaluated at the entire sample average contribution of $4,273 per year, children from a household with a parent with a college degree have a 15 percent probability of no college education versus 37 percent for those from a household without a parent with college degree, a difference of 22 percentage points. The gap is even wider when factoring in the fact that more educated households contribute more money toward their children's education. For example, the probability of a child having no college education is only 9 percent for households with a college-educated parent contributing $6,336 per year (the average for higher-educated households). The probability of no college education for a child from a household without a college-educated parent is 45 percent when contributing $2,054, the average for such households.

The parameters estimates also indicate a sizeable impact from parental contributions to the educational attainment of their children. For example, an additional dollar reduces the probability of no college education for a child by as much as 0.007 percentage points to 0.001 percentage points for college-educated households. So that means an additional $1,000 contributed per year will lower the probability of a child getting no college education by 1 to 7 percent. The impact is the greatest when they are not contributing, and then steadily falls, which is partly because the probability of no college falls so low. This marginal effect is smaller at first for households without a college-educated parent, but remains more stable as more is offered, though additional dollars still have a decreasing marginal impact. At the same time, an additional dollar offered increases the probability of a four-year degree, and this marginal impact rises some before eventually declining as well. For example, the marginal impact on the probability of a four-year degree for college-educated households begins at 0.005 for the first dollar, rises to over 0.0065 and then falls back below 0.005. Again a similar pattern is true for less educated households, but the marginal impact is smaller more stable.

A few other parameter estimates are worth mentioning at this point. The discount rate for a six-year period, δ, is estimated at 0.8797 which is the equivalent of 0.979 per year. This is actually less discounting than would be implied by the real interest rate, which is set at three percent. The child-cost estimates are $77,228 per six years for college-educated households and $41,782 for those without a college degree, both large amounts for the respective households to cover. The estimate for γ, the coefficient of relative risk aversion, is 2.8. While prior estimates of relative risk aversion range quite a bit, this is in line with typical estimates. Lastly, the estimated taste parameters for number of children and for children's education, while not of interest in size

---

[23] The marginal effect of variable, $x_k \in x$, on the probability that certain choice j is selected is given

themselves, are suggestive in their relation between college-educated households and less-educated households. The increase in utility from an additional child is greater for less-educated households, and does not decline as rapidly with increasing numbers of children. The utility of additional education is also surprisingly greater for less-educated households, but in this case the decrease in marginal utility is more rapid for those households. This may be necessary to get enough college spending from those households, while college-educated households both have more money and the probability of their children getting a college degree is much higher.

5.2 Model Fit

Before turning to experiments on educational savings policies, briefly consider some information on the fit of the model. Table 1 presents some summary statistics for the actual data and as predicted by a simulated sample of 10,000 households based on the model parameter estimates. As the table shows, the model does well in matching most average characteristics in the data. The model solidly captures the increasing accumulation of assets, though it slightly understates average assets levels for younger households, overstates them in middle age and understates them again as households move into their fifties. This is true for both college-educated and less-educated households. Still, the averages are not too far off in general. Comparing figure 3 with figure 1 shows that the model clearly captures the qualitative trends seen for the actual data for savings patterns and the number of children in a household.

The model does a reasonably good job in matching the percentage of children with different levels of education. For example, for households with a college-educated parent the model overstates the percentage with no college education by just 0.9 percent, understates the percentage with some college by just 0.56 percent and understates the percentage with a 4 year degree by just 0.34 percent. For households without a parent with a college degree, the predicted percentages are equally good. For all households together, however, the percentages are off by a bit more, largely because the model does a slightly less good job with predicting the relative number of children in households of different education levels.

The model simulation matches average amounts contributed by households to their children's college education quite well. As table 7 shows, overall the simulated sample average

---

by: $\frac{\partial \Pr\text{ob}(\text{outcome } j \mid x)}{\partial x_k} = [\phi(\mu_{j-1} - x'\kappa) - \phi(\mu_j - x'\kappa)]\frac{\partial x'\kappa}{\partial x_k}$

is $4,364 per-child, per-year, while the actual data average is just slightly less at $4,273. The mean comparison is equally good when looking within different education levels. However, some distributional aspects of these offers are clearly not matched as well. The model noticeably underpredicts the percentage of parents actually spending money on their children's education. The model predicts 70.69 percent of households contributing to their children's education, while the actual data show 81.71 percent contributing. This discrepancy shows up in both household education levels: 61.87 percent in the model simulation versus 67.72 percent in the data for less-educated households, and 82.80 percent versus 94.71 percent for households with a college degree. These lower percentages contributing are accompanied by a wider dispersion of offers, particularly on the high end, within the simulated data versus the actual data.

When looking at the number of children, the model again does fairly well in predicting the average number of children, but again misses a bit more on some aspects of the broader distribution. The average number of children in the data is 2.31 versus 2.28 in the simulated data[24]. Beyond just the overall average, the model also does well in matching the average number of children born at different ages for a household, and when looking separately ad households of different education levels.

6. Education Policy Simulations

The section here uses the model parameter estimates and a simulation of 10,000 households, proportioned by education as is in the data, to examine how their savings and educational contribution decisions would change, and the impact of these changes on children's education attainment, when different policies are introduced. In particular, I will consider the creation of tax-advantaged educational savings accounts. For a comparison, I will also look at the introduction of a direct college subsidy and a tax credit.[25] More specifically, the educational savings account is allowed to grow tax-free and the earnings are not taxable as long as they are spent on education.[26] Any amount not spent on education becomes taxable. Since there are no

---

[24] Remember that the data here are limited to married families with a stable marriage, so these are higher than the total population.

[25] While the tax reform act of 1997 did introduce and/or expand policies such as these, assuming that the policies were largely unexpected by households, they were not impacted by them over their life-cycle paths.

[26] For this experiment, instead of just one type of assets, there are now two types of savings: regular savings and educational savings. The addition of the second asset class does create an additional state variable in the model and it is a second continuous variable. This creates a much larger state space across which the problem must be solved,

taxes directly in the model, this tax-free savings is modeled as an increase in the return on savings equal to the amount of the default return times the parents' marginal tax rate. For tax rates, I use the federal 2001 tax brackets for the marginal rates on married households filing jointly.[27] Parents can only contribute to these accounts if they have children and cannot contribute more than $7,500 per child per year. When the child turns 18, which is period 4 in the model, the amount in the account must either be contributed toward education (which is obviously contingent on the child attending college) or a penalty is paid. The penalty is an amount equal to the current marginal tax rate times any unspent balance in the account. The subsidy is a flat $1,000 education grant. This appears in the model as a default offer to all children, so that, for example, when parents offer nothing, the child still has a $1000 offer. Finally, I consider a tax credit available on all money spent on children's education. In the model this basically means that parents don't actually have to pay out part of their offered contribution when children go to school, and the savings is equal to their marginal tax rate times the contribution. Table 5 presents the results of the policy experiments. The reported amounts are the sample statistics for the appropriate simulated data set. The base simulation is the original model as estimated and discussed above.

The education savings account actually gives the largest increase in contributions from parents. Over all households, the average parental contribution increases $2,072, to $6,436, from $4364 in the base simulation. Surprisingly, the increase is actually greater for households where neither parent has a college degree. For college-educated parents the increase is only $1549, while for those without a college degree the increase is $2,233. While this result is most likely driven by the fact that the marginal impact of additional contributions is decreasing and more educated households already contribute more, it is not what is generally expected of a policy that is often presumed to benefit wealthier households more. Both groups of households also see an increase the percentage of parents contributing: from 82.80 percent to 90.21 percent for parents with a college degree and from 61.87 percent to 74.10 percent for parents without a college

---

thought the extended model does not need to be estimated, just get an obtainable solution. The same solution procedure is used as before, except that now the Keane-Wolpin smoothing regressions are now applied across both asset classes

[27] These are 15% for income less than $45,200, 28% for income between $45,200 and $109,250, 31% for income between $109,250 and $166,500, 36% for income between $166,500 and $297,350 and 39.6% for income above $297,350. Since the income in the model is over six year periods, I divide this by six to get the associated tax bracket.

degree. With greater contributions and an overall parental contribution rate of 82.56 percent, the percentage of children with no college education fall by 6 percent, to 35 percent from 41 percent, and the percentage with a four-year degree rises to 27.92 percent from 23.23 percent. The cost of this program, in the form of lost revenue from taxable earnings on the education savings, comes to almost $8,000 per household. However, much of this lost revenue is on earning from savings that were not present from the base scenario, making this estimate hard to compare. If the amount of savings is held to the base simulated levels, the lost revenue is just $2,593 per household. The difference is because the education savings accounts greatly alters savings behavior. Average savings increases by 20 percent by ages 24-29, 28 percent by ages 30-35 and still by 16 percent for ages 46-41. After that, assets levels are largely unchanged for ages 42-47 and then actually one percent lower for ages 48-53. The average amounts in the education accounts are also shown in table 8 and show that households do actually increase savings and not just shift assets.

For a comparison point of these results, I also look at a grand and tax credit. These policies are much more in line with those examined in previous studies described in section 2, though here again I am considering more directly the impact on parents behavior. The $1,000 grant has the smallest impact on the average amount contributed by parents, and the average is lower than in the base simulation. However, that is to be expected since every child is in effect already receiving a $1,000 per year contribution, and the marginal impact of contributions is decreasing as contributions increase. Interestingly, though, the average contribution from parents does not decline by a full $1,000, so on net more is being contributed in total toward education. For all households, the average contribution falls by $681 so the net amount from parents and the grant increases by $319. For households with a college degree there is an average net increase of $427, and for parents without a college degree, the average net increase is $200. The percentage of parents contributing also falls noticeably, but again 100 percent of children are receiving a $1,000 contribution from the new grant. So while only 55 percent of parents contribute above and beyond the grant, a decline of just over 15 percentage points from the percentage contributing before the grant, there is still a sizeable impact on the distribution of education outcomes for children. The percentage of children with no college education falls by almost five percent, from 40.6 percent to 35.8 percent. This is consistent with the 4% and 4.5% attendance increases estimated by Dynarski (2002) and Ichimura and Tabor (2002), respectively, and gives additional validation of the model. The percentage with a college degree rises by 3.3 percent from 23.23 percent to 26.54 percent. Furthermore, like the ESA, the improvement is actually greater for children whose parents have no college education, even though the grant itself was not means tested and the net increase from parents was smaller. For these children, the percentage receiving no college education actually falls by 6.3 percent, from 52.4 percent to 46.1 percent. The improved educational attainment, while larger in total with the ESA, is larger for less-educated households with the grant, and the expected cost of the grant comes to $4,146 per household, between the range of estimated cost of the ESA above.

The tax credit generates a larger increase in average contributions than the grant, but still less than the ESA. It also increases the percentage of parents contributing, but the net impact on educational outcomes is less. The average contribution rises to $5,963, an increase of almost $1,600. This is true for households of both education levels, those with a college degree show a $1,509 increase and those without a college degree show a $1,681 increase. Furthermore the percentage of parents contributing rises from 70.69 percent to 82.55 percent. The increase is the greatest among parents without a college degree, where the percentage contributing rises from 61.87 percent to 80.16 percent, while the increase is just 3 percentage points, from 82.8 percent to 85.8 percent, for parents with a college degree. This suggests that while the credit did increase the amount of support from college educated parents who gave support, it did not greatly increase the percentage contributing. However, the impact on educational outcomes is the lowest here. The percentage receiving no college degree does fall, but only by 2.6 percent (from 40.59 percent to 37.99 percent), which is noticeably less than the 4.8 percent decline with the grant and 5.7 percent with the ESA. Despite this, the expected cost of the credit is $5,775 per household, more than the grant. Lastly, and interestingly, the grant and tax credit do not actually have much impact at all on parent savings.

7. Conclusion

The model presented here gives structural estimates of a dynamic, life-cycle model with the choices of having children, spending on children's college education, and savings and consumption all made endogenously. The model also allows for borrowing constraints, uncertain lifetime income, and allows for heterogeneity between parents with different levels of education. The model is solved using backward induction and a regression smoothing procedure and is then estimated with a simulated maximum likelihood procedure using data from the National Longitudinal Survey. The estimated model is used to run policy experiments to gauge the impact of several programs aiming to increase parental support for children's education.

The model estimates that the vast majority of patents, over 70 percent, offer to aid in college expenses and that the average amount is significant, at over $4,000 per year, per child. Furthermore, the size of the offer has a sizeable impact on the educational attainment of children. For example, an additional $1,000 will reduce the probability a child receives no education by anywhere from 4 percent to 7 percent. The marginal impact of this additional money is decreasing as the offer size increases, and the impact differs for parents with a college degree and those without a college degree. Furthermore, the estimates indicate that children's educational attainment is also directly influenced by the education level of parents.

The policy experiments examine the impact of an education grant/subsidy, tax breaks for education spending, and the creation of a targeted savings account. The experiments suggest that a $1,000 grant available to all students and the creation of tax-free savings accounts with a limit of $7,500 a year both have a greater impact on education outcomes than the tax credit. The grant is effective because it in essence increases the percentage of households offering support to 100%, and households do not decrease the amounts they offer beyond this by the full $1,000. Furthermore, given that the initial money offered has the largest marginal impact, getting offer rates up is very beneficial. The savings account, on the other hand, vastly raises the amount that parents offer. This is because the return on using the accounts is higher. Furthermore, the percentage of households making offers also rises substantially as the benefit of having children with post-secondary education becomes more attainable. Interestingly, the greatest increase in both contributions and improving probabilities of receiving a college education are among lower educated households, even with the education savings account, which is not necessarily what is expected.

A number of limitation in the model used here still remain. The education decisions of children are not modeled in a fundamental way, and future work to include this decision more directly, perhaps worked into a cooperative or noncooperative framework with parents would allow for a better understanding of this process. Futhermore, there is no allowance in the model for studying the choice of college to attend. Other studies present evidence that there is substantial substitution, for example, between private and public four-year schools in response to price increases (see for example, Cameron and Heckman, 1999.) In addition, there are no differences in the quality of college education here, though it is very likely that parental contributions may influence the "quality" of school attended. Lastly, there is no direct accounting for the negative impact that greater parental assets and contributions might have on reducing financial aid that would otherwise be available. In the estimation here, such effects are just evident in the falling marginal impact of parental contributions. Allowing for these factors within a dynamic framework, and finding data rich enough to estimate the such a multi-faceted analysis, is clearly a challenge for future research.

References

Acton, F. S. (1990). Numerical Methods That Work. Washington, D.C., Mathematical Association of America.

Attanasio, O. P. and M. Browning (1995). "Consumption over the Life Cycle and over the Business Cycle." American Economic Review **85**(5): 1118-37.

Berkner, L., A. Berker, et al. (2002). Student Financing of Undergraduate Education: 1999-2000. Washington, D.C., U.S.Department of Education, National Center for Education Statistics.

Bowman, A. W. (1985). "A Comparative Study of Some Kernel-Based Nonparametric Density Estimators." Journal of Statistical Computation and Simulation **21**(3-4): 313-327.

Browning, M. and A. Lusardi (1996). "Household Saving: Micro Theories and Micro Facts." Journal of Economic Literature **34**(4): 1797-1855.

Caballero, R. J. (1991). "Earnings Uncertainty and Aggregate Wealth Accumulation." American Economic Review **81**(4): 859-71.

Choy, S. P. and R. R. Henke (1992). Parental Financial Support for Undergraduate Education. Washington, D.C., U.S. Department of Education, National Center for Education Statistics.

Coleman, A. (1998). Household Savings: A Survey of Recent Microeconomic Theory and Evidence. New Zealand Treasury Working Paper.

Deaton, A. (1991). "Saving and Liquidity Constraints." Econometrica **59**(5): 1221-48.

Dynarski, S. (2000). "Hope for Whom? Financial Aid for the Middle Class and Its Impact on College Attendance." National Tax Journal **53**(3): 629-61.

Dynarski, S. (2002). "The Behavioral and Distributional Implications of Aid for College." American Economic Review **92**(2): 279-285.

Geweke, J., M. P. Keane, et al. (1994). "Alternative Computational Approaches to Inference in the Multinomial Probit Model." Review of Economics and Statistics **76**(4): 609-32.

Härdle, W. (1991). Smoothing Techniques with Implementation in S. New York, Springer-Verlag.

Haveman, R. and B. Wolfe (1995). "The Determinants of Children's Attainments: A Review of Methods and Findings." Journal of Economic Literature **33**(4): 1829-78.

Hubbard, R. G. and K. L. Judd (1987). "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax." American Economic Review **77**(4): 630-46.

Hubbard, R. G., J. Skinner, et al. (1994). "Expanding the Life-Cycle Model: Precautionary Saving and Public Policy." American Economic Review **84**(2): 174-79.

Hurd, M. D. (1989). "Mortality Risk and Bequests." Econometrica **57**(4): 779-813.

Ichimura, H. and C. Taber (2002). "Semiparametric Reduced-Form Estimation of Tuition Subsidies." American Economic Review **92**(2): 286-292.

Keane, M. P. (2002). "Financial Aid, Borrowing Constraints, and College Attendance: Evidence from Structural Estimates." American Economic Review **92**(2): 293-297.

Keane, M. P. and K. I. Wolpin (1994). "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence." Review of Economics and Statistics **76**(4): 648-72.

Keane, M. P. and K. I. Wolpin (1997). "The Career Decisions of Young Men." Journal of Political Economy **105**(3): 473-522.

Keane, M. P. and K. I. Wolpin (2001). "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment." International Economic Review **42**(4): 1051-1103.

Kotlikoff, L. J. (1989). What determines savings? Cambridge, Mass. and London, MIT Press.

Kronholz, J. (2003). College Education Becomes Costlier. The Wall Street Journal. New York**:** 3.

McFadden, D. (1989). "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration." Econometrica **57**(5): 995-1026.

Scott, D. W. (1992). Multivariate Density Estimation: Theory, Practice, and Visualization. New York, John Wiley and Sons.

Silverman, B. W. (1986). Density Estimation for Statistics and Data Analysis. New York, Chapman and Hall.

Skinner, J. (1988). "Risky Income, Life Cycle Consumption, and Precautionary Savings." Journal of Monetary Economics **22**(2): 237-55.

Stern, S. (1994). "Two Dynamic Discrete Choice Estimation Problems and Simulation Method Solutions." Review of Economics and Statistics **76**(4): 695-702.

Titterington, D. M. and A. W. Bowman (1985). "A Comparative Study of Smoothing Procedures for Ordered Catagorical Data." Journal of Statistical Computation and Simulation **21**(3-4): 291-312.

Zeldes, S. P. (1989). "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." Quarterly Journal of Economics **104**(2): 275-98.

TABLE 1

ACTUAL AND PREDICTED SELECTED OUTCOMES

	Actual	Predicted
All Households:		
Assets* (mean by age):		
24-29	24,140	22,292
30-35	59,858	57,925
36-41	85,820	86,193
42-47	131,628	138,729
48-53	218,947	215,224
Children (mean)	2.313	2.282
Children born when 18-23 (mean)	0.896	0.887
Children born when 24-29 (mean)	0.890	0.876
Children born when 30-35 (mean)	0.527	0.520
Contributions to College* (mean)	4,273	4,364
Percent Contributing	81.71	70.69
Childrens Education Attainment (percent):		
No College	42.37	40.59
Some College	35.73	36.18
4-year degree or more	21.90	23.23
By Education:		
Households with High School Degree or Less:		
Assets* (mean by age):		
24-29	19,291	18,643
30-35	45,943	43,968
36-41	62,472	59,811
42-47	89,729	98,315
48-53	157,730	155,317
Children (mean)	2.395	2.370
Children born when 18-23 (mean)	1.225	1.235
Children born when 24-29 (mean)	0.826	0.804
Children born when 30-35 (mean)	0.344	0.331
Contributions to College* (mean)	2,054	2,185
Percent Contributing	67.72	61.87
Childrens Education Attainment (percent):		
No College	52.49	52.43
Some College	33.72	34.23
4-year degree or more	13.79	13.34

* 1999 dollars

continued…

TABLE 1

CONTINUED

	Actual	Predicted
Households with College Degree:		
Assets* (mean by age):		
24-29	30,296	25,654
30-35	77,523	75,642
36-41	115,458	119,681
42-47	184,814	190,031
48-53	304,100	300,348
Children (mean)	2.208	2.117
Children born when 18-23 (mean)	0.478	0.445
Children born when 24-29 (mean)	0.971	0.967
Children born when 30-35 (mean)	0.759	0.759
Contributions to College* (mean)	6,336	6,390
Percent Contributing	94.71	82.80
Childrens Education Attainment (percent):		
No College	23.10	24.00
Some College	39.56	39.00
4-year degree or more	37.34	37.00

* 1999 dollars

TABLE 2
NET ASSETS*

Age (No. Obs.)	Mean (Median)	Std. Dev.	Min.	Max.	Percent Negative	Mean if Negative
All						
18-23	4,290	10,330	-9,480	77,443	10.07	-1,172
(556)	(965)					
24-29	24,140	36,968	-12,776	297,912	8.81	-1,597
(556)	(12,342)					
30-35	59,858	74,375	-20,561	587,955	6.29	-2,685
(556)	(41,957)					
36-41	85,820	104,145	-28,095	673,160	7.01	-3,399
(556)	(56,331)					
42-47	131,628	148,201	-11,266	1,093,175	3.60	-1,965
(556)	(85,364)					
48-53	218,947	245,800	-18,950	1,196,500	1.89	-3,026
(318)	(133,470)					
54-59	263,326	298,600	-5,000	1,246,000	2.44	-1,620
(41)	(145,000)					
By Parents Education: High School or Less						
18-23	4,431	11,028	-8,617	77,443	9.32	-933
(311)	(287)					
24-29	19,291	30,695	-10,029	196,752	10.61	-1,591
(311)	(5,628)					
30-35	45,943	61,429	-16,477	401,446	8.04	-2,391
(311)	(28,166)					
36-41	62,472	86,877	-28,095	661,441	10.29	-3,202
(311)	(38,024)					
42-47	89,729	111,792	-11,266	790,741	6.11	-1,908
(311)	(57,938)					
48-53	157,730	199,906	-10,000	1,126,000	1.62	-1,493
(185)	(85,000)					
54-59	174,348	237,751	-5,000	945,500	6.67	-1,620
(30)	(82,500)					

* 1999 dollars

continued…

TABLE 2

CONTINUED

Age (No. Obs.)	Mean (Median)	Std. Dev.	Min.	Max.	Percent Negative	Mean if Negative
By Parents Education:						
College						
18-23	4,111	9,391	-9,479	64,789	11.02	-1,657
(245)	(1,555)					
24-29	30,296	42,935	-12,776	297,913	6.53	-1,611
(245)	(17,887)					
30-35	77,523	85,017	-20,561	587,955	4.08	-3,673
(245)	(56,912)					
36-41	115,458	116,209	-27,462	673,160	2.86	-4,495
(245)	(80,958)					
42-47	184,814	170,301	-8,416	1,093,175	0.41	-2,872
(245)	(137,546)					
48-53	304,100	277,183	-18,950	1,196,500	2.26	-6,091
(133)	(213,700)					
54-59	505,996	322,794	144,000	1,246,000	-	-
(11)	(456,000)					

TABLE 3

PARAMETER ESTIMATES

Utility Function

γ	λ_{1c}	λ_{11c}	λ_{12c}	λ_{13c}	λ_{1h}	λ_{11h}	λ_{12h}	λ_{13h}
2.8263	0.8271[a]	5.3555[a]	5.4439[a]	6.0019[a]	-0.2912[a]	5.9955[a]	8.6061[a]	1.7350[a]
(0.0250)	(0.5402)	(0.1457)	(1.4300)	(1.6707)	(0.0767)	(0.0388)	(1.7221)	(0.2395)
λ_{2c}	λ_{2h}	λ_3	α_{1c}	α_{1h}	α_{2c}	α_{2h}	θ_{1c}	θ_{1h}
-0.6286[a]	0.1431[a]	0.0011[b]	0.0041[a]	0.0053[a]	-0.3699[b]	-0.5472[b]	0.0087[a]	0.0122[a]
(0.5526)	(0.5514)	(0.0054)	(0.0000)	(0.0001)	(0.4773)	(1.7105)	(0.0001)	(0.0004)
θ_{2c}	θ_{2h}							
-0.5421[b]	-0.8246[b]							
(0.3131)	(2.9127)							

Income

β_{c0}	β_{c1}	β_{c2}	β_{c3}	β_{h0}	β_{h1}	β_{h2}	β_{h3}
11.0244	0.3348	-0.0354	-0.1927	0.1684	-0.0065	0.9626	-0.2011
(0.2387)	(0.0401)	(0.0017)	(0.0992)	(0.2111)	(0.0322)	(0.0005)	(0.0834)
b_c	b_h						
0.9575	0.0088						
(0.2235)	(0.3240)						

Children's Education

μ_1	μ_2	κ_1	κ_2	κ_3	κ_4	κ_5
0.07189	1.2974	1.1023[c]	-3.8649[a]	0.4040	0.7243[c]	0.1011[a]
(0.0110)	(0.0153)	(0.0874)	(1.5150)	(0.0157)	(0.0616)	(0.6696)

Error Distribution

σ_c	σ_n	σ_e	σ_{ic}	σ_{ih}	σ_{retc}	σ_{reth}	σ_η
0.0101	7.9767[a]	1.3856[c]	0.7082	0.5182	42275	23275	0.0285
(0.0076)	(0.5700)	(0.9334)	(0.0129)	(0.0060)	(21395)	(3799)	(0.1411)

Other Parameters

δ	ψ_c	ψ_n
0.8797	77228	41782
(0.0147)	(5736)	(3814)

[a]Parameter multiplied by 10^8

[b]Parameter multiplied by 10^{11}

[c]Parameter multiplied by 1000

TABLE 4

ESTIMATED EDUCATION PROBABILITIES BY PARENT CONTIBUTION AND EDUCATION

	Parental Contribution				
	$0	$2,054	$4,273	$6,336	$10,000
Households with College Degree:					
Prob(no college)	36.99%	24.47%	14.83%	9.04%	3.74%
Marginal impact of $1	-0.000069	-0.000055	-0.000039	-0.000026	-0.000012
Prob(some college)	44.43%	45.87%	42.38%	36.47%	25.15%
Marginal impact of $1	0.000020	-0.000006	-0.000027	-0.000037	-0.000038
Prob(4-year degree)	18.58%	27.66%	42.80%	54.49%	71.12%
Marginal impact of $1	0.000049	0.000061	0.000065	0.000063	0.000050
High School Degree or Less:					
Prob(no college)	52.87%	44.50%	37.12%	31.87%	25.98%
Marginal impact of $1	-0.000044	-0.000040	-0.000035	-0.000031	-0.000023
Prob(some college)	37.41%	41.65%	44.39%	45.59%	45.98%
Marginal impact of $1	0.000025	0.000018	0.000010	0.000005	-0.000001
Prob(4-year degree)	9.73%	13.85%	18.49%	22.54%	28.04%
Marginal impact of $1	0.000019	0.000023	0.000025	0.000026	0.000024

TABLE 5
POLICY EXPERIMENTS

	Base	Ed. Sav. Account	Grant	Tax Credit
All Households:				
Parent Contributions to College (mean)	4,364	6,436	3,605	5,963
Percent of Parents Contributing	70.69	82.56	54.51	82.55
Childrens Education (percent):				
No College	40.59	34.91	35.80	37.99
Some College	36.18	37.17	37.66	36.65
4-year degree or more	23.23	27.92	26.54	25.36
Children (mean)	2.282	2.410	2.285	2.283
Total Assets (mean by age):				
24-29	22,292	26,812	22,286	22,295
30-35	57,925	73,905	57,914	57,941
36-41	86,193	100,009	86,148	86,259
42-47	138,729	139,286	139,246	138,245
48-53	219,224	216,791	221,448	218,627
Educational Savings (mean by age):				
24-29		1,360		
30-35		13,759		
36-41		31,525		
42-47		51,321		
48-53		45,756		
By Household Education:				
High School Degree or Less:				
Parent Contributions to College (mean)	2,185	4,418	1,385	3,866
Percent of Parents Contributing	61.87	74.10	43.72	80.16
Childrens Education (percent):				
No College	52.43	48.03	46.10	49.29
Some College	34.23	35.91	37.13	35.30
4-year degree or more	13.34	16.06	16.77	15.41
Children (mean)	2.370	2.644	2.373	2.371
Total Assets (mean by age):				
24-29	19,643	29,878	19,647	19,647
30-35	43,968	57,222	43,964	43,990
36-41	59,811	76,654	59,760	59,921
42-47	98,315	102,638	98,861	97,553
48-53	155,317	144,641	156,567	153,905
Educational Savings (mean by age):				
24-29		1,185		
30-35		11,267		
36-41		26,002		
42-47		46,814		
48-53		42,370		

continued…

TABLE 5
CONTINUED

	Base	Ed. Sav. Account	Grant	Tax Credit
Households with College Degree:				
Parent Contributions to College (mean)	6,390	7,939	5,817	7,899
Percent of Parents Contributing	82.80	90.21	69.32	85.84
Childrens Education(percent):				
No College	24.00	20.16	21.48	22.31
Some College	39.00	39.20	38.40	38.51
4-year degree or more	37.00	40.64	40.12	39.18
Children (mean)	2.117	2.223	2.118	2.117
Total Assets (mean by age):				
24-29	25,654	27,011	25,656	25,656
30-35	75,642	94,099	75,624	75,650
36-41	119,681	128,734	119,644	119,692
42-47	190,031	195,500	190,511	189,898
48-53	300,348	280,819	303,807	300,783
Educational Savings (mean by age):				
24-29		1,816		
30-35		18,531		
36-41		38,532		
42-47		57,041		
48-53		50,168		

FIGURE 1
ASSET ACCUMULATION BY NUMBER OF CHILDREN

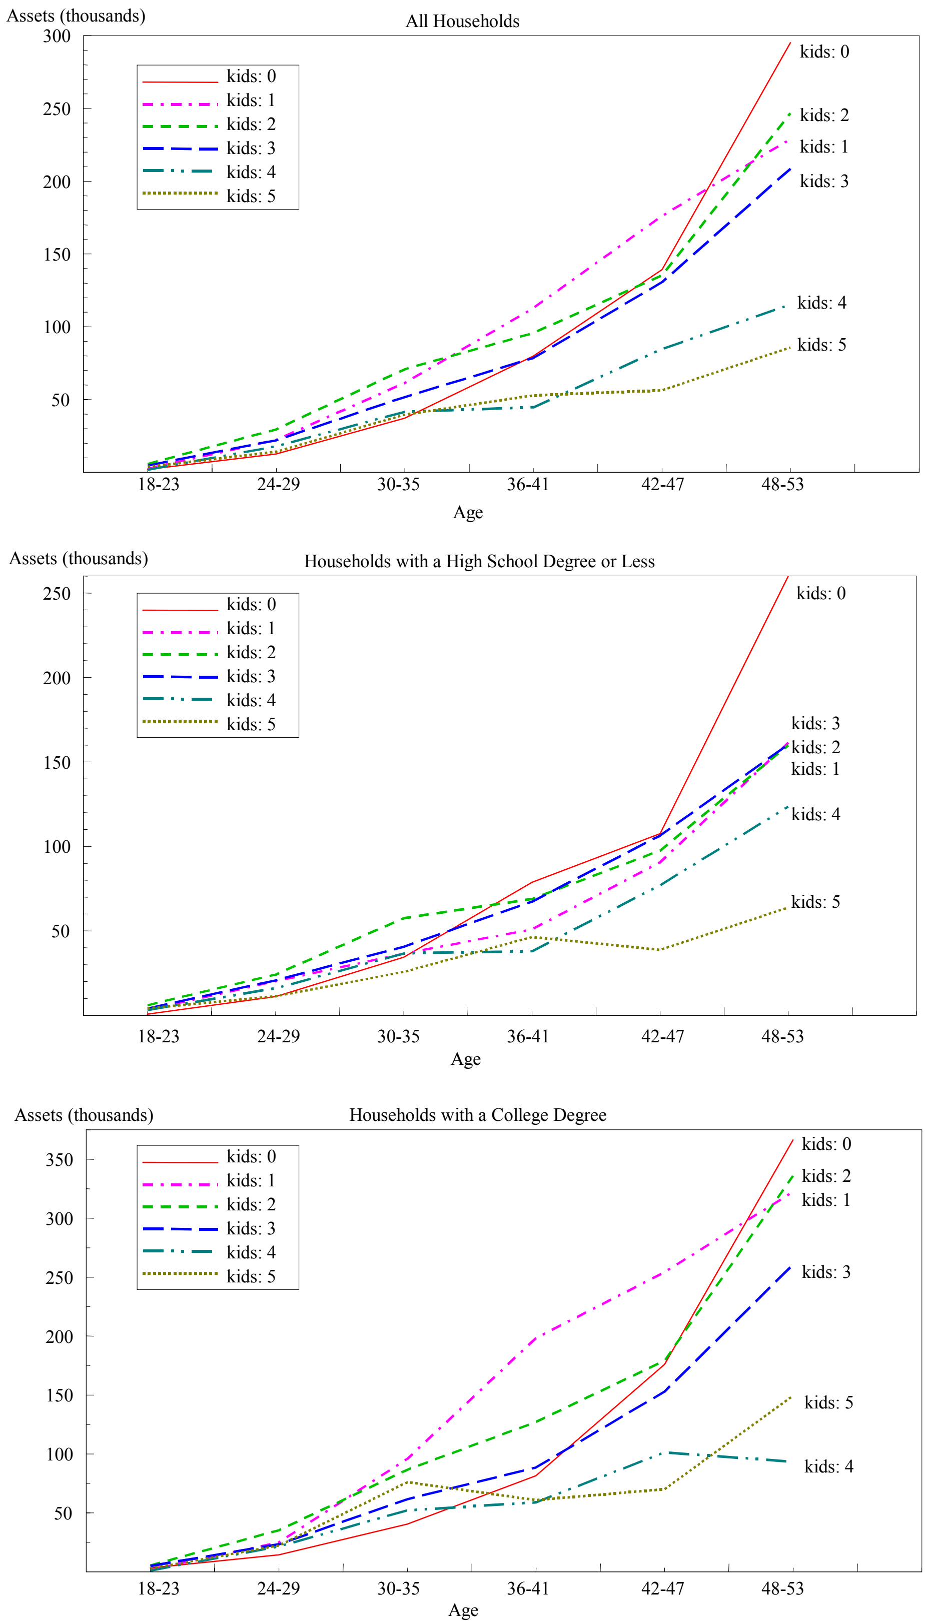

FIGURE 2. Exact Functional Forms

Utility Function:

$$U_t = \frac{(c_t \varepsilon_{ct})^{1-\gamma}}{1-\gamma} + (\lambda_1 + \lambda_2 n_t + \varepsilon_{nt})(n_t + \lambda_3 c_t) + (\alpha_1 + \alpha_2 q_{mt} + \varepsilon_{qt}) q_{mt} + (\theta_1 + \theta_2 q_{ht} + \varepsilon_{qt}) q_{ht}$$

where in addition:

$$\begin{aligned}
\lambda_1 = {} & \lambda_{1c} + \lambda_{11c} 1(ed \geq college, t=1) + \lambda_{12c} 1(ed \geq college, t=2) + \\
& \lambda_{13c} 1(ed \geq college, t=3) + \lambda_{1h} 1(ed \leq high\ school) + \\
& \lambda_{11h} 1(ed \leq high\ school, t=1) + \lambda_{12h} 1(ed \leq high\ school, t=2) + \\
& \lambda_{13h} 1(ed \leq high\ school, t=3)
\end{aligned}$$

$$\lambda_2 = \lambda_{2c} + \lambda_{2h} 1(ed \leq high\ school)$$

$$\alpha_1 = \alpha_{1c} + \alpha_{1h} 1(ed \leq high\ school)$$

$$\alpha_2 = \alpha_{2c} + \alpha_{2h} 1(ed \leq high\ school)$$

$$\theta_1 = \theta_{1c} + \theta_{1h} 1(ed \leq high\ school)$$

$$\theta_2 = \theta_{12} + \theta_{2h} 1(ed \leq high\ school)$$

Income Function:

Before Retirement: $I_{jt} = \exp[\beta_{j0} + \beta_{j1} t + \beta_{j2} t^2 + \beta_{j3} 1(t=1) + \varepsilon_{ijt}]$

After Retirement: $I_{jt} = b_j E\left[\frac{1}{T^{**}-1} \sum_{t=1}^{T^{**}-1} I_{jt}\right] + \varepsilon_{rjt}$

For $j \in \{\geq college, \leq high\ school)$

Education Attainment Function:

$$S_t = \kappa_1 O_t + \kappa_2 O_t^2 + \kappa_3 1(ed \geq college) + \kappa_4 O_t 1(ed \geq college) + \kappa_5 O_t^2 1(ed \geq college)$$

$$\Pr(no\ college) = \Phi(\mu_1 - S_t)$$

$$\Pr(some\ college) = \Phi(\mu_2 - S_t) - \Phi(\mu_1 - S_t)$$

$$\Pr(college\ \text{degree}) = 1 - \Phi(\mu_2 - S_t)$$

Error Distributions:

$\ln(\varepsilon_{ct}) \sim N(0, \sigma_c^2)$, $\varepsilon_{nt} \sim N(0, \sigma_n^2)$, $\varepsilon_{qt} \sim N(0, \sigma_q^2)$, $\varepsilon_{ijt} \sim N(0, \sigma_{ij}^2)$, $\varepsilon_{rjt} \sim N(0, \sigma_{rj}^2)$

Measurement Error:

$assets_{it}^{obs} = assets_{it}^{true} \exp(\varepsilon_\eta)$ with $\varepsilon_\eta \sim N(0, \sigma_\eta^2)$

Other Parameters:

Discount Factor: δ

Child Costs: Ψ_j for $j \in \{\geq college, \leq high\ school)$

FIGURE 3
SIMULATED ASSET ACCUMULATION BY NUMBER OF CHILDREN

Assets (thousands) All Households

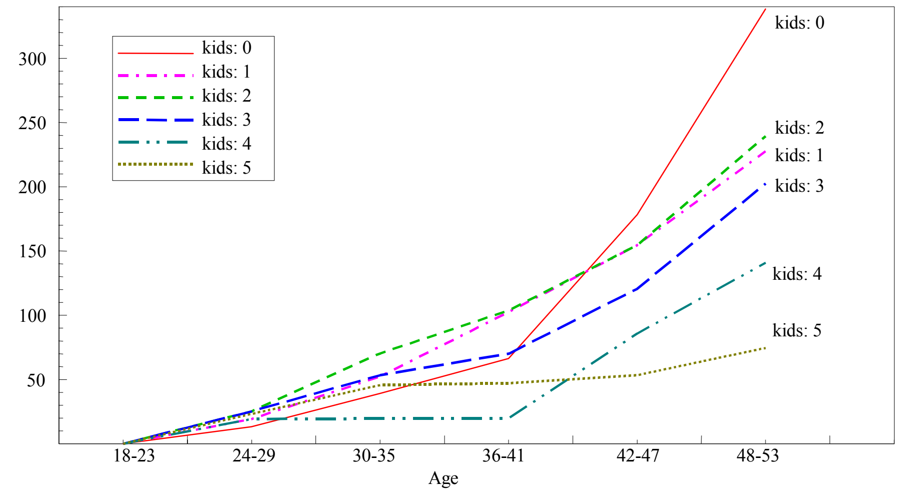

Assets (thousands) Households with a High School Degree or Less

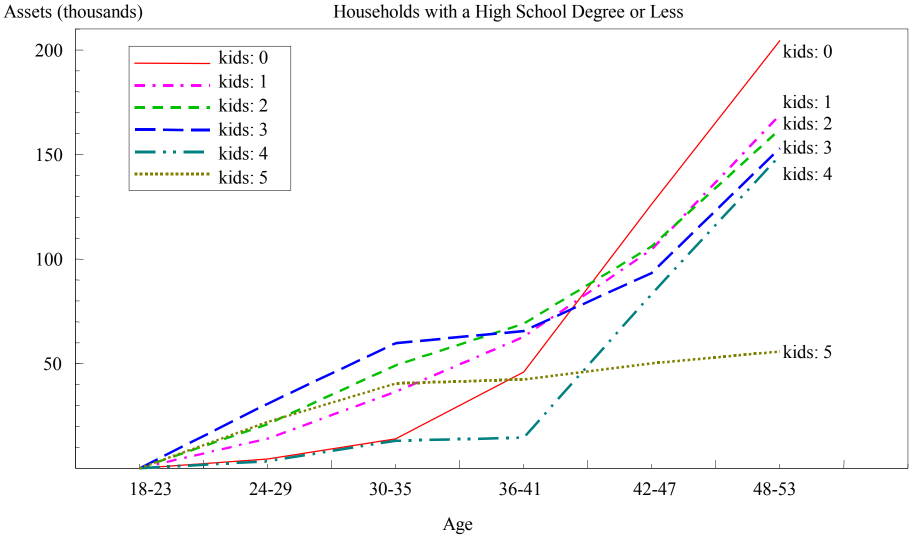

Assets (thousands) Households with a College Degree