1-1-2003

The impact of grants, tax credit and education savings account on parental contributions to college expenses and the educational attainment of children

Michael D. Morris
University of New Orleans

THE IMPACT OF GRANTS, TAX CREDITS AND EDUCATION SAVINGS ACCOUNTS ON PARENTAL CONTRIBUTIONS TO COLLEGE EXPENSES AND THE EDUCATIONAL ATTAINMENT OF CHILDREN

Michael D. S. Morris*

DRAFT: February 12, 2003

*Assistant Professor of Economics, University of New Orleans, Department of Economics and Finance, 2000 Lakeshore Drive, New Orleans, LA 70148; phone: (504)-280-6900; fax: (504)-280-6397; email: mdmorris@uno.edu

THE IMPACT OF GRANTS, TAX CREDITS AND EDUCATION SAVINGS ACCOUNTS ON PARENTAL CONTRIBUTIONS TO COLLEGE EXPENSES AND THE EDUCATIONAL ATTAINMENT OF CHILDREN

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 the dynamic programming model solution to create simulated data samples from which nonparametric kernel estimators are used to construct the densities in 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 life-cycle paths of these choices suggest that parents do save to help make sizeable transfers to their children, and that making such choices endogenous is important. Furthermore, 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 additional government grants for college education, tax credits for college spending and the creation of tax-free education savings accounts on parental savings, contributions toward education, and the education attainment of children. While all of the policies increase net contributions to children and increase the probability that a child attains a college degree, the grants and education savings accounts are found to be the most effective. In addition, both policies are actually found to have a greater impact on children with less educated parents.

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 cost of a college degree is also very large. 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 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 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 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 in recent years a push to help parents pay for their children's college education. In the last five years, there has been an 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. Such parental contributions toward education can have differing impacts on savings. First, it may increase savings early in life if parents plan to spend money on their children's education. Working in the other direction, however, is the negative impact of the additional cost of raising children. Lastly, at the time of the transfer and beyond, assets may change their rate of growth. To fully examine the impact of policies on parental contributions, then, it is necessary to consider parents 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 borrowing constraints, assumes that households save to smooth consumption. 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. Taken together, there are several motives for households to save: precautionary, liquidity constrained, consumption smoothing for retirement and the provision of inter-vivos transfers to children in the form of paying 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 policy experiments aimed to replicate an increase in grants for children attending college, an increase in tax credit for parental spending on education, and the creation of tax-free college savings accounts.

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 model and estimation procedure in detail. A general life-cycle model is developed with the addition of choices on having children and helping to pay for their education and a specification is presented that will be estimated here. The section also discusses the issues associated with solving the model and the numerical technique used. The section concludes with a description of the simulated maximum likelihood procedure that is 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 aimed at capturing the impact of grants, tax credits and tax-free education savings accounts on both parental contributions and the education attainment of their children.

2. Background and Previous Literature

The workhorse for economic research on household intertemporal 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). The life-cycle framework takes its original inspirations from Modigliani and Brumberg (1954), but has evolved significantly through the years. The basic underlying tenet of these models is that forward-looking households will try to equate their marginal utilities across different periods by smoothing consumption across their lifetime. The original framework progressed into what is widely referred to as the Certainty-Equivalence (CEQ) model that dominated much of the research from the mid 1970's through the late 1980's.

The validity of the CEQ model has been debated and tested for years. Generally speaking, there are several empirical shortcomings from the predictions of the CEQ-model (for summaries see Hubbard, Skinner et al. 1994; Browning and Lusardi 1996; Coleman 1998). The additions of precautionary savings and liquidity constraints to the life-cycle framework have been both successful and rather widely accepted improvements in life-cycle models (see for example and review Hubbard, Skinner et al. 1994; Hubbard, Skinner et al. 1995). These models are complicated by the fact that they can in general no longer be analytically solved and must therefore rely on numerical solution techniques. Several sources of uncertainty confronting households have been studied in these precautionary models. Early studies of precautionary savings allow for earnings risk (Skinner 1988; Zeldes 1989; Caballero 1991; Deaton 1991), lifetime uncertainty (Hubbard and Judd 1987; Hurd 1989) and even uncertainty regarding medical expenses (Kotlikoff 1989). The general finding in these studies is that the precautionary and liquidity constrained models do a better job of accounting for the micro facts in the data than the CEQ model. The borrowing constraint prevents households from borrowing to consume more, even when that would be optimal, and this allows for an increased correlation between income and consumption. Furthermore, the precautionary motive encourages households to save a "buffer stock".

While there is strong evidence that households do face liquidity constraints, the previous discussion of the debate on this illustrates, there are other choices that households make that are related to savings and consumtion decisions that are ignored by life-cycle model. 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. In particular, findings that demographic equivalency scales can help reconcile model predictions to the data, as noted by Attanasio and Browning (1995), indicate the importance of family size in considering the evidence. However, as Keane and Wolpin point out, such scales are "fundamentally arbitrary" and just indicate that standard models can fit the data fairly well "by allowing for enough interactions between consumptions and household demographics in the utility function (Keane and Wolpin 2001, p. 1058)." In fact, a major shortcoming of life-cycle models when considering family size is that the number and timing of children is either taken as given or as the result of a random, exogenous process. Coleman claims the lack of allowing for "the endogeneity of the overall relationship between family structure, family income, family consumptions and age," is one of the "greatest weakness[es] in the current state of theory (Coleman 1998, p. 8)." Furthermore, Browning and Lusardi add that it "would be extremely interesting to extend these studies … to allow for the fact that, to a certain extent, agents choose … the time path of demographics. Thus high education agents marry later and start their families later which may be connected to their income processes (Browning and Lusardi 1996, pp. 1807-8)."

There can be no denying that the consumption and savings behavior of families will differ depending on the number of children in the family. For starters, children undoubtedly cost money. Using data from the 1990-1992 Consumer Expenditure Survey (CES) and adjusting the amounts to 1997 dollars using the CPI, Lino (1998) calculates the average annual expense for the younger child, under age 18, in a two-child household for different household income brackets. He calculates the expense ranges from $5,820 to $6,880 depending on the age of the child (the amount increases as children grow) for households with 1997 pre-tax income less than $35,500[1]. This average increases to a

[1] The costs include data on child specific expenses on clothing, child care and education and estimated portions of expenses on housing, transportation, food, health care and miscellaneous expenditures.

range of $8,060 to $9,170 for households with income between $35,500 and $59,700 and to $11,990 to $13,260 for households with income greater than $59,700. These amounts average 14% to 28% of a household's pre-tax income depending on the pre-tax income bracket of the household. These explicit costs, of course, do not include any additional opportunity costs a household may incur from labor choices made in response to raising children.

While both the intended or accidental transfers to children in the form of bequests and whether households plan on receiving these bequest may be debatable, inter-vivos transfers (transfers from parents to children while the parents are still living) must be intended and the plans for such transfers are much more likely to be considered by the recipients when making decisions. These transfers often occur as gifts, paying for college, or as help with a down payment on a house. Using data from the 1983-1986 Survey of Consumer Finances, Gale and Scholz (1994) find that intended inter-vivos transfers can make up approximately 20% of a household's wealth. However, again, these are often concentrated among the very wealthy. Keane and Wolpin (2001) estimate a model where young households receive transfers from parents (the process is random but allowed to depend on several know factors) and find that inter-vivos wealth transfers received by young households are significant. To the extent that such transfers do account for wealth accumulation, they should be considered when modeling life-cycle behavior, both for the parents and the children. Furthermore, these decisions are certainly related to the choice of having children and represent a very complicated challenge to model.

Transfers from parents to children often take the form of investments in human capital. In particular, it is very common in the U.S. for parents to help pay for the post-secondary education of their children. These educational subsidies do not have some of the divisive issues involved with the other wealth transfers discussed above. Helping pay for college, like other inter-vivos transfers, is certainly intended by parents and not accidental. Furthermore, since they are used toward paying the expenses of receiving an education, they are not so much direct additions to the asset wealth of children as they are an education subsidy and determinant in the future income path of children.

One reason for the prevalence of these transfers is the rising out-of-pocket expense of college education in the U.S. Berkner et al. (2002), using data from the National Center for Education Statistics (NCES), report that for that the average tuition for private, not-for-profit, 4-year schools was $15,000 for the 1999-2000 academic year. This was lower at other schools: $8,900 for private, for-profit institutions, $4,300 for public 4-year institutions and $1,600 for public 2-year institutions. The full price including non-tuition expenses is much higher. The same study finds that for a full time student, the average full price was $23,600 for private, not-for-profit, 4-year schools, $18,400 for private, for-profit schools, $12,600 for public, 4-year schools and $9,100 for public 2-year schools. These are obviously significant and large expenses.

Of course, these costs are not necessarily the out-of-pocket expense incurred by students because many receive some assistance from grants, loans, and/or other forms of financial aid. Berkner et al. (2002) find that for the 1999-2000 academic year, 27% of students received some for of aid but not loans, 22% received subsidized loans and other aid, and 7% received subsidized loans only. That still leaves 45% with no aid at all. The percent will no aid falls to 28% when considering only full-time students, but that is still a significant portion without aid. More importantly, the aid sources do not come close to covering the full cost of education for most students. Berkner et al. (2002) calculate that the average total aid for full time students receiving some form of aid at private, not-for-profit 4-year institutions in 1999-2000 was $11,600, and was $7,200 for students at private, for profit schools, $6,200 for public 4-year schools and $2,300 for public 2-year schools. These findings are consistent with McPherson and Shapiro (1991) who a decade earlier found that 38% of students in public schools and 68% of those in private schools received some form of aid but that on average it only covered 75% of public school expenses and 50% of private school expenses.

The bottom line is that students still must cover large out-of-pocket expenses for post-secondary education. Some of the money needed certainly comes from students' own earnings and savings. In addition, many students receive financial support from their parents, and these amounts can be large. Using data from the National Postsecondary Student Aid Study (NPSAS), Lee (2001) finds that for the 1995-1996 academic year, 91.9% of undergraduate students attending schools with tuition and fees greater than $12,000 received some direct financial contribution from their parents. This percentage was still high for students at schools with tuition and fees less than $12,000: 79.6% for public research institutions and 70.8% for other institutions. Using data gathered by the NCES, Choy and Henke (1992) found 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). While the percentage of parents contributing and the amounts they contribute have surely been increasing as college costs have risen, these contributions have long been a significant source of funds for students. Leslie (1984), using data from the Cooperative Institutional Research Program (CIRP), found that for the 1979-1980 academic year 47.6% of students received support from their parents and the amount averaged $1,426 ($3,272 in 1999 dollars).

The out-of-pocket expenses that students must pay for a college education has also motivated a wide variety of policies to help students pay afford a college degree. Most government support in the form of the grants and loans previously discussed is directly targeted at the student attending school. Several studies have tried to measure the extent these programs help increase college attendance. For example, many argue that an education subsidy should have a significant impact, especially if students tend to face borrowing constraints. 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%. She also claims that this is consistent with previous findings. Ichimura and Taber (2002) estimate a 4.5% increase in attendance from the availability of a $1,000 subsidy. They estimate this using a reduced form estimation derived from the model used by Keane and Wolpin (2001) and use the same 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, Keane and Wolpin, while finding that borrowing constraints are indeed tight for students (they can not even support one year from borrowing alone), they 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 significantly increase the educational attainment of children, which is consistent with the findings that subsidies increase attendance.

In recent years, several policies have been created to help parents pay for the college education of their children. That is, policies where the goal is to increase the amounts parents contribute toward their children's education. The Taxpayer Relief Act of 1997 created several such policies. The first two are tax breaks. 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. 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). The credit can be claimed for expenses for yourself or a dependant. Hence, parents can claim it for college expenses paid for their dependant children. The Lifetime Learning Credit is similar except that it is available for any year of education (not just the first two years) but is for 20% of up to $5,000 in expenses (this will increase to $10,000 in 2003). Also, the Hope Scholarship Credit can be claimed for each child supported separately but the Lifetime Learning credit is a per family credit. The act also set up an educational savings account, originally often referred to as an educational IRA. This is really just a custodial account for children under 18 for which the earnings and distributions are tax-free as long as they are used for qualified education expenses. However, 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 and the program was renamed the Coverdell Education Savings Account. An additional type of savings plan that was created are 529 plans (also created in 1997, they are named for the IRS tax code that allows for favorable tax treatment to qualified state tuition programs) which are actually run by states. Since they are administered by states, the details can vary some from state to state but in general there are two types of 529 plans. The first type is a prepaid tuition plan. In these plans, 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. 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. As such, these plans are really just hedging against rapid rises in college tuition. The other type of 529 plan is a college savings plan. These accounts accept contributions that can grow tax-free until distribution. Starting in 2002, withdraws were also allowed to be tax-free. Specific limits vary by states, but at the federal level all contributions are treated under the gift tax laws, which are not triggered until gifts exceed $10,000. Starting in 2002, families were also allowed to switch between different types of 529 plans. More recently, President Bush's budget proposal for 2003 includes the creation new education savings accounts that would replace the Coverdell accounts and allow for $7,500 in contributions per year.

The model and work presented here builds upon the above strains of literature in several important ways. First, 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. The educational subsidies in the model will also provide a way to conduct policy simulations to look at predicted impacts of programs aimed at getting parents to contribute more toward their children's education.

3. **The Model**

3.1 The General Model

This section develops a general model to look at household consumption and savings decisions along with choices on how many children to have and how much to spend on children's college education. A key feature of the model, unlike previous models looking at household savings interacted with the number of children in the household, is that the number of children in a household is an endogenous choice. The model is a based on a household 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 how much to save. In addition, younger households must also decide how many children to have, and households with college-age children must decide how much money to contribute to their children's education.

To make these decisions, households maximize their expected discounted utility. Each household's contemporaneous utility for a period depends on both the level of consumption and the number of children in the family. Furthermore, households receive 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} is a consumption taste shock at time *t*, ε_{nt} is a taste shock to the number of children at time *t*, and ε_{et} is a taste shock to the education levels at time *t*.

In each period, households choose a level of consumption and savings. Households can also have children in periods 1 through T^*, and so must also choose how many additional children to have in each of those periods. As stated previously, children may add to a household's utility both directly and through the amount of education they attain. There are also, however, costs to having children. While children live at home, the household incurs a cost of Ψ per child. All children are assumed to live at home for *M* periods, at which time they move out.

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 observably demographic variables. For the household, letting d_{jt} be the education level child *j* receives in period *t*,

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

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

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

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 as mentioned previously. The expectation is over the taste shocks to consumption, children, their educational attainment and the realizations of future income and children's educational attainment. To implement this model, then, what remains is 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. These details, along with a method to both solve the model and to estimate the model parameters are laid out in the next sections.

3.2 Model Specification

I now turn to a full specification of the general model. The specification will allow the model to be solved and subsequently estimated using the NLS data described in section four. In particular, specifications are needed for the utility function, the random taste shocks, the income and children's educational attainment processes, the parental offer process, 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.

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[2]. 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 in any one period, and no more than 5 children in total. These restrictions impact only a small portion of the data and significantly ease the

---

[2] 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.

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).

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), the household has the option to help pay for that child's college education. It is assumed that households only make these contributions in the period that their children move out. 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.

When making an offer of financial support for college education, households 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 offers make on how much schooling children receive. This is necessary because what are reported in the data are the amounts parents actually contributed toward their children's education. Obviously, then, 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. The data suggest that this is not overly restrictive. Again when looking at the NLS data, 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 that on average, when comparing families with the same number of children, the families with more children in college will be making larger contributions toward 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. The data, though, do support the assumption and its implication. Looking at the NLS data, for families with two children born 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.

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.[3]

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

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

which combines a constant relative risk aversion (CRRA) utility function for consumption with additively separate quadratic utility in the number of children, 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 shocks for each period are drawn as follows. In every period, $\varepsilon_{ct,}$ is drawn from an iid lognormal distribution with log mean zero and ε_{nt} and ε_{qt} are drawn from zero mean iid normal distributions, so that

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

Household income is assumed to be determined by,

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

where the characteristic vector Z_t contains dummies for the levels of education and marital status interacted with a quadratic in age. The exponential specification for earnings leads to a standard wage equation where $ln(w_t) = Z_t\beta + \varepsilon_{It}$ for pre-retirement income. 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 by zero mean iid normal distributions so

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

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 decisions to marry and divorce and the separate contributions toward their children's education from divorced couples would also need to be modeled. Assuming a constant marital status allows for a simpler model that focuses on the household savings and educational transfer decisions. This clearly will limit the sample this model is applied to, which will be outlined in more detail later. The level of

---

[3] In addition, limiting the number of outcomes to three will greatly ease the computation burden of solving

education of the household for earnings consideration is also a constant in this model. As such, both marital status and education are assumed to be a constant family type.

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

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

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, a household is not allowed to offer more than can be covered without taking out an uncollateralized loan.

The timing of the model is outlined in figures 1-3. 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, the household then receives their income and pays out expenses for children living at home. The household then makes 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 the agent has children who are 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 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. 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

$$\begin{aligned} &O_t(s) = \max_o E_d\left[V_t(s') \mid o, s\right] && \text{if have children to make offer to} \\ &O_t(s) = V_t(s) && \text{otherwise} \end{aligned} \tag{17}$$

subject to (5)-(16), 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

$$\begin{aligned} &V_t(s) = \max_{c,k',n'}\left\{U(c_t, n_t, q_{mt}, q_{ht}, \varepsilon_{ct}, \varepsilon_{nt}, \varepsilon_{ht}) + \delta E\left[O_{t+1}(s') \mid s\right]\right\} && \text{if t = 3,4,5} \\ &V_t(s) = \max_{c,k',n'}\left\{U(c_t, n_t, q_{mt}, q_{ht}, \varepsilon_{ct}, \varepsilon_{nt}, \varepsilon_{ht}) + \delta E\left[V_{t+1}(s') \mid s\right]\right\} && \text{otherwise} \end{aligned} \tag{18}$$

subject to (5)-(16) 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_l or ε_r. It is assumed that $V_{T+1}(s) = 0$. In most periods, then, the household just goes from solving the primary problem in one period to solving it again in the next period. After periods 3, 4, and 5, however, the household first solves the offer problem and then proceeds to solve the primary problem for that period.

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 making their decisions, households actually use the level of disposable cash-on-hand, and not the asset level directly. Let x_t be the amount of disposable cash-on-hand. At the beginning of every period,

$$x_t = (1+r_k)k_{t-1} + I_t - n_{ht}\psi. \tag{19}$$

Within the periods that the offer and savings problems are both solved, x_t updates by subtracting the amount that was spent on college.

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. 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. 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 (17) and (18).)

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. As outlined previously, these restrictions on the choice set only impact a small portion of the data, and they clearly ease the computation of a model solution. 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 such, the model here clearly still suffers from a curse of dimensionality.

One method of overcoming this is to utilize an estimate for the integration involved in the expectations and to solve the model for a subset of the possible state points and interpolate for the remaining values of the EMAX. Keane and Wolpin (1994) propose using Monte Carlo integration to evaluate the expectations in the EMAX and interpolating over a subset of state points for which the model is solved using least squares. 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 (18) the interpolation is the regression of directly solved EMAX values on the contemporary utility evaluated at the means of the stochastic components and for (17) the optimal offers are regressed on assets (Keane and Wolpin, 1994).

3.4 Estimation

The model is estimated using Simulated Maximum Likelihood[4]. 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 (Stern 1994; Keane and Wolpin 1997; 2001). 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

---

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

accomplish this with an innovative use of assuming there is measurement error for the unobserved variables. With discrete choice models the likelihood is composed of the probabilities of the different choices occurring. 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). Since the model in this paper contains continuous choices, the exact same procedure cannot be followed. However, the same concept applies, but instead of probability simulators I use nonparametric a density estimator is used with simulated samples to construct the elements of the likelihood.

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

$$L_i = \prod_{t=0}^{T-1} f(s_{t+1}^i \mid s_t^i) \tag{20}$$

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.

A problem remains in calculating the likelihood because the functional fomr of $f(.)$ is unknown. However, given a set of the parameters, 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 $f(s_{t+1}|s_t)$. Several smooth density estimators have been proposed and used for discrete choice models in a setting such as this (e.g. McFadden 1989; Geweke, Keane et al. 1994; Stern 1994; Keane and Wolpin 1997). These methods consist of using some sort of smoothed probability estimator of the probabilities of the different finite (discrete) outcomes occurring[5]. However, the model here involves a continuous choice in assets. One possible way to deal with this is to discretize assets. Yet, the model and solution methods have been constructed to preserve a continuous choice in assets. An alternative, then, is to use a continuous density estimator. This is the approach that I take in estimating the model.

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. In using a continuous variable, though, the density is no longer the probabililty of one outcome resulting from a finite number of choices. Instead, the continuous or mixed density of the next period state space must be estimated. Furthermore, the dimension of the state space becomes important. Unfortunately, traditional kernel estimators are notoriously innacuarate and difficult (if not impossible) to implement when used in higher dimension space. Since the state space here is certainly multidimensional, this must be addressed.

Given 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 estimated density must be calculated. The remaining values are fixed. As such, it is never necessary to estimate a joint density for more than two variables. For example, after period six, the number of children and their education levels are all fixed so just assets change. As such, the conditional density for these periods is just the one-dimensional conditional density of next period assets given current assets. 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

[5] Several of the previous estimates of discrete choice models referenced above have constructed their density estimates based on the fact that a certain set of observed outcomes (state space values) is the result of the value function evaluated at those outcomes being the greatest among the finite (dicrete) options available. As such, the probabilities can be calculated from simulating the probability the value function is greatest for these choices.

offer and the initial state, and finally of savings given the initial state and educational outcomes. So 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 being decided upon. As such, again only a bivariate density estimator is needed for these periods.

The specific density estimators that I use are as follows. In general, a univariate kernel density estimate is calculated from

$$\hat{f}(s) = \frac{1}{nh}\sum_{i=1}^{n} K(s, S_i, h) \tag{21}$$

where *s* is the data point at which the density is to be calculated, *n* is the number of data points in the data sample, S_i is the ith observation in the data sample and h is a smoothing parameter. The function *K(.)* is the kernel function and the resulting $\hat{f}$ is the estimate for the true density, *f*. When a univariate continuous density estimator is needed, in periods seven and beyond (for assets) and in the offer portion of periods four through six (for the offer value), I use a standard Gaussian Kernel. As such the kernel function is

$$K(s, S_i, h) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{s - S_i}{h}\right)^2\right) \tag{22}$$

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

$$h = 1.06\hat{\sigma}n^{-1/5} \tag{23}$$

where $\hat{\sigma}$ is the standard deviation of the variable in the sample data. This choice of a value for the smoothing parameter (or 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.

For the bivariate densities I use a bivariate product kernel (Scott 1992) which in general takes the form

$$\hat{f}(s_1, s_2) = \frac{1}{nh_1h_2}\sum_{i=1}^{n} K_1(s_1, S_{i1}, h_1)K_2(s_2, S_{i2}, h_2). \tag{24}$$

When the bivariate estimates are needed, however, 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 (18)-(19). For the discrete values, a variety of smooth kernel estimators could be used. Notice, though, that the variables here (number of children and number of children with certain educational outcomes) are not just categorical, but are ordered and the values matter. For example, having three children is closer to the choice of having two children than say having five children. 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

$$K(s, S_i, h) = \lambda^{|s - S_i|} \tag{25}$$

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

$$h = \sum_{j=0}^{J-1} \lambda^j \tag{26}$$

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 is set to 0.3[6].

So 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 routing using a version of a simplex algorithm originally proposed by Nelder and Mead (1965). and discussed in Acton (1990).

4. The Data

The data are taken from the young women cohort of the National Longitudinal Survey (NLS). The young women cohort of the NLS 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 original cohort sample is based on a core random sample and an over-sample of blacks. While the NLS collects data on a wide variety of topics, the relevant information used in this analysis is data on assets, income, children, education, marital status, children's education and spending on children's education. The data used here are from 1968-1999.

4.1 Sample and variable definitions.

Periods in the model are specified to be six years long. When matching this 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.

[6] 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

Comprehensive questions on assets were asked in 1968, 1971-1973, 1978, 1983, 1988, 1993, 1995, 1997 and 1999. Asset information obtained includes information about housing, mortgages, savings, stocks, bonds, vehicles, business, farm or real estate assets and outstanding debts. For analysis with the model above, assets are measured as total net household assets, excluding vehicles. If more than one observation of assets is available within the appropriate age range for a period for a household (which is the case when assets information is collected regularly), 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. Income questions 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 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. 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 collected in every survey. Information on the highest grade attended is collected and updated every survey. In addition, in 1983 questions were asked on whether the respondent had earned a high school diploma or GED and if so when. This information was then updated in all subsequent surveys for those respondents who attended school since the last interview. Finally, information on the highest degree received was collected in the initial year and updated in every subsequent interview for those who had received a degree since the last interview. Information on the level of education of spouses is also collected. The level of education enters the model in two ways. First, it is a factor in the wage equation. Second, it is a factor in the level of education attained by children in the household. Also, the model assumes that the level of education is constant for a household's lifetime.

---

a reasonable range does not qualitatively impact the estimates. For more discussion on kernel methods for discrete distributions refer to Aitken (1983).

The household education level is therefore 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 one of two categories, those with a college degree, and those without. So, there are in essence 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, and because the estimation technique takes multiplicatively longer for each additional education level.

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, in this cohort, 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. As such, the results in actuality can only be claimed to be attributed to such stable and married households.

In every interview, detailed information is collected for each of the respondent's children living in the household. This information includes among other things, the child's age. In addition, in 1978 information on the age of up to eight children born by 1978 was collected. Similar information followed in 1983 for children born between 1978 and 1983. In 1985, 1988, 1991, 1993, 1995 and 1999 the information was collected for children born since the last interview. 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 was explained in detailing the model specification, 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. As was mentioned when discussing these restrictions in the specification, 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 the parent's provided toward college for each child within the past twelve months. The offer in the model is assumed to be the same for all children born in the same period, as was described earlier along with the reasons for this assumption and appropriateness of the assumption. 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 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.

As mentioned previously, 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.[7] This leaves a sample of 556 households with 3,139 household-period observations.

4.2 Descriptive Statistics

Table 1 shows the sample breakdown between of the average number of children for different levels of household education. 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. During ages 30-35, college educated households have an average of 0.759 children, while those without a college degree have only 0.344.

Table 2 presents 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. In fact, the accumulation is very rapid from an average of just over $4,000 in 1999 dollars at ages 18-23 to an average of $218,947 for when reaching ages 48-53. The amount of assets is greater for those with a college degree and the accumulation even more rapid. Still, for those with no more than a high school degree, average assets grow from $4,431 when 18-23 to over $170,000 by the time they reach ages 54-59. As usual, the asset distribution is largely skewed toward high values as is shown by the much lower median asset values. For the entire sample, the median asset level grows from a mere 956 when aged 18-23 to $145,000 when reaching ages 54-59.

The percent of households with negative asset values is also listed in Table 2. Ten percent of the sample lists negative asset levels at ages 18-23 but this progressively declines to only 3.6 percent when aged 42-47 and 1.89 percent by ages 48-53. The percentage declines less steadily through the years for households without a college

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

degree, but still is never greater than just over ten percent. 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 percentage of households do have negative net assets, this 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 4 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 accumulate assets more slowly in early years, but then at a much more rapid pace so that by their late forties, households with no children have more assets on average. In fact, at young ages households with no children have the fewest assets, 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 4. 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.

While this is certainly optional for parents to help pay for a college education, most households do contribute toward their children's education and the amounts are significant. Table 3 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. This data comes from parents who answered affirmatively in the surveys that they had a child in college within the last 12 months, and if so how much they contributed toward that child's education. Given the biennial collection of the data, this allows for some gaps in off years when children might attend college some, but no information is colleted on how much parents helped to pay. However, that should not significantly change the results in Table 3.

Table 4 shows 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). Furthermore, as noted above, parents with a college degree make larger contributions toward helping pay for a college education, which should also increase number of children with a college degree (as is found to be the case in Keane and Wolpin 2001).

5. Results

5.1 Parameter Estimates

The estimated parameters of the model are reported in table 5. There are a total of 48 parameters estimated here. While that is a large number of parameters, the data consist of over 3,139 household-period observations. These are then fit to the various choices in the model including a choice of zero to four children in each of the first three periods, an amount to contribute toward education for the fourth period of children's lives, and an amount to save in each period. Moreover, the latter two are allowed to be continuous choice variables.

Several of the estimated parameters are of direct interest by themselves. In the specified CRRA utility function for consumption, γ is the coefficient of relative risk aversion, $1/\gamma$ is the intertemporal elasticity of substitution and $\gamma+1$ is the coefficient of relative prudence (Kimball, 1990). The estimate for γ, the coefficient of relative risk aversion, is 2.8. This in turn implies an intertemporal elasticity of substitution of 0.357 and a coefficient of prudence is 3.8. While prior estimates relative risk aversion and elasticity of substitution vary quite a bit, these estimates are in line with typical estimates. Hubbard et al. (1994) discuss previous estimates when deciding on what value to use in their analysis and conclude that a typical estimate for the coefficient of relative risk aversion is 3, which is what they proceed to use. The estimate here, while indicating slightly more willingness to make intertemporal substitutions, is very similar. Looking at other studies indicates that if anything the estimate here of the willingness to make intertemporal substitution is on the low end. Hurd (1989) estimates intertemporal substitution of between 0.89 and 1.4. Keane and Wolpin (2001) estimate a coefficient of relative risk aversion of 0.5 which implies a much lower coefficient of relative prudence of 1.5. They attribute this to their specific limits on borrowing constraints keeping those with steep earnings profiles, but still facing earnings uncertainty, from borrowing heavily as opposed to relying on a high degree of prudence. The estimate here suggests more prudence even with borrowing constraints. The difference is presumably a result of the modeling of the endogenous fertility choice and that the desire to afford children and spending on education provide a possible explanation for the greater prudence. The decision to have a child locks the household into some additional expenses for three periods, and the spending on education is a one-period chance that the household cannot intertemporally substitute into other periods. Both of these factors would seem to encourage a greater degree of prudence for households.

The estimated parameters for the children's educational outcome suggest that both parental offers for financing and the parents' education level both 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.[8] Table 6 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. Note that these predictions are drawn just from the ordered probit and not model simulation predictions, which allow for the endogenous choice of the offers. For evaluating the impact of the level of parents' education, the marginal impact can be seen by comparing the change in the predicted probabilities. However, again the marginal effect is contingent on the contribution amount. 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 probabilities of a four-year college degree are 43 percent versus 18 percent, an increase in 25 percentage points for children with a college-educated parent. The gap is even wider when factoring in the fact that more educated households also 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. This is a gap of 36 percentage points. Evaluated at these offer amounts, the probabilities of a child with a four-year degree are 54 percent versus 14 percent, a difference of 40 percentage points.

The parameters estimates also indicate a sizeable impact from parental contributions to the educational attainment of their children. The marginal impact of additional parental contributions changes as the amount contributed changes. The marginal impact also differs between households with a parent with a college degree and those without. 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

---

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

year will lower the probability of a child getting no college education by 1 to 7 percent. The impact is the greatest for college-educated households when they are not contributing, and then steadily falls for these households, 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. Still, additional dollars have a decreasing marginal impact. For example, the marginal effect of the first dollar offered reduces the probability of no college by 0.004 percentage points and this slows reduces to 0.002 as parents offer more. At the same time, an additional dollar offered increases the probability of a four-year degree, but this marginal impact rises and then eventually falls 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, and hence more stable, rising from 0.0019, up to 0.0026 and then falling slightly. The marginal impact on the probability of some college education starts positive and then eventually becomes negative as the probability of a four-year degree grows. So, in summary, the parameter estimates indicate a sizeable impact of both parents' education and the offer to help pay for college on the probability a child receives a college education. Furthermore, this impact is magnified by the fact that the marginal impact of an additional dollar is often larger for educated households, at least at low spending levels, though the impact diminishes as more is spent for all households.

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. Lastly, the estimated taste parameters for number of children and for children's education, while not of interest in size

---

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

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 a further discussion of the implications of the model on educational savings policies and the fertility decisions in subsequent sections, I now turn to providing some information on the fit of the model. Table 7 presents some summary statistics for the actual data and 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. In general the model solidly captures the increasing accumulation of assets. The model does slightly understate the average assets levels for younger households, overstate them in middle age and understate them again as households move into their fifties. This is true for both college-educated and less-educated households as is shown graphically in figure 5. The averages are not too far off in general. At most the means are off by 15 percent, for college educated households when they ages 24-29 and 10 percent for households without a college degree when they are ages 42-47. The remainder of the time the average is off by less than 10 percent.

Additional information on the fit of predicted asset levels is provided in figures 6-8. Figure 6 shows the predicted asset accumulation for households broken down by both education levels and number of children. The model clearly captures the qualitative trends seen for the actual data in figure 4. Households with no kids save less in early years, but have the highest asset levels by their fifties. Households with children save more early on, but have lower average asset levels by their fifties, especially those with four or five children. Figure 7 shows separate asset accumulation panels for households with different numbers of children, comparing the paths of the actual data and the simulated data. Again the qualitative trends are clearly captured, including the changing pace of asset accumulation. However, especially for households with 4 and 5 children, the actual levels are over or understated by a noticeable amount. However, the sample size for these households is definitely smaller. Figure 8 plot the cumulative distribution of assets for households in different age brackets and different education levels. As seen here, when looks at all households, the general trend is for the predicted data to be more disperse than the actual data. That is the predicted distribution has fatter tails.

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, as will be mentioned in more detail below. Still, the simulated sample only understates the percentage with no college degree by 1.78 percent, overstates the percentage with some college by .45 percent and overstates the percentage with a four-year degree by 1.33 percent.

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 is $4,364 per-child, per-year, while the actual data average is just slightly less at $4,273. For households with a college degree the average is higher, at $6,390 for the simulated data and $6,336 for the actual data, but again very little difference. The same is true for less educated households where the average in the simulated sample is $2,185 versus $2,054 in the actual data. 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[9]. Beyond just the overall average, the model also does well in matching the average number of children born at different ages for a household. The averages for mothers aged 18-23 years old, 24-29 years old and 30-35 years old are 0.896, 0.890 and 0.527 respectively in the actual data. The model simulation matches this quite will with averages of 0.887, 0.876 and 0.520. The fit is similarly good when looking separately ad households of different education levels, as is also shown in table 7.

6. Education Policy Simulations

Given the model estimates, policy experiments can be made to evaluate the impact of policies that would impact the amount that parents contribute to their children's education or the impact that these contributions have. As was discussed in section 2, , there are a variety of educational reforms aimed at making a college education more affordable and/or helping parents save for their children's education. The section here uses 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 in the presence of policies similar to an increase direct college subsidies, a tax credit, a tax credit with an income limit and the introduction of a tax-advantaged educational savings account. The change in the distribution of education levels of children is also examined. While the tax reform act of 1997 did introduce and/or expand policies such as these, assuming that the

---

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

policies were largely unexpected by households, the households used to estimate the model certainly were not impacted by them over their life-cycle paths.

The specific policy adjustments that are considered are as follows. The first is a \$1,000 education grant. This appears in the model as a default offer to all children, that is the amount when parents offer nothing, of \$1,000 instead of nothing. The second policy is a tax credit to parents for all money spent on their 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. This is complicated by the fact that the original model embeds the changing tax code and tax brackets throughout the lives of the households. In the experiment I just use the federal 2001 tax brackets for the marginal rates on married households filing jointly.[10] As such, the credit is really like a partial rebate on your college expenditures that depends on the amount you contribute and your current income. The third policy is this same credit, but with a cap at \$40,000 (\$10,000 for four years) on the amount of contributions eligible for the credit. The final policy is the addition of an education savings account. For this experiment, instead of just one type of assets, there are now two types of savings: regular savings and educational savings. The educational savings account is allowed to grow tax-free and the earnings are not taxable as long as they are spent on education. Any amount not spent on education becomes taxable. Since there are no 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. 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 of the parents times any unspent balance in the account. 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

[10] 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

a much larger state space across which the problem must be solved, 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.

Table 10 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 $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 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 the average contribution by parents declines by $573, indicating an increase in total contributions including the grant of $427. Similarly for parents without a college degree, there is decline of parental contributions of $800, which means an increase in total contribution of $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. The percentage with some college education rises by 1.5 percent, from 36.2 percent to 37.7 percent and the percentage with a college degree rises by 3.3 percent from 23.23 percent to 26.54 percent. Furthermore, the improvement is actually greater for children whose parents have no college education, even though the grant itself was not means tested. For these children, the percentage receiving no college education actually falls by 6.3 percent, from 52.4 percent to 46.1 percent.

---

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.

The tax credit generates a larger increase in average contributions than the grant and 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. Certainly this is partly due to the already high percentage of educated parents who give support.

While the amounts spent by parents did go up on average, and the increase was greater than with the grant experiment, the impact on educational outcomes was less. The percentage receiving no college degree did 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. The associated increase in the percentages with some college and a four-year degree are0 .47 percentage points and 2.13 percentage points respectively. Compared to the grant, these increases are smaller, though the relative shift toward a four-year degree is more. This is presumably due to the fact that the grant is available to all children, which in effect ups the contribution percentage to 100 percent. Coupling this with the fact that the estimated marginal impact of the offer is generally decreasing, this can explain the superior performance of the grant. Interestingly, the simulated expected per-household cost of the grant is lower than the lost tax revenue from the credit. The expected cost of the grant comes to $4,146 per household, while the expected credit costs are $5,775 per household. Imposing the high, $10,000 per year, cap on the amount eligible for the credit does reduce this cost to $5,364. The cap also lowers the average contribution by almost $300, when compared to having the credit without the cap in place. However, the cap made no significant difference to the percentage of parents contributing, nor on the distribution of educational outcomes. For example, the cap only reduced the decline in the percentage of children with no college education by 0.16 percentage points. This is because with the cap so high, it is not a factor in getting people to contribute, and it only reduces contributions for those families who already contributed so much that the marginal impact of the declining offers is not large.

One interesting result from these first three experiments is that there is virtually no impact on the number of children that families have or on their asset accumulation paths at younger ages. The average number of children iss exactly the same before and after the tax credit experiments, and rose only by 0.001 in the grant experiment. Furthermore, when looking at asset accumulation of parents up to age 41, there is not much difference at all from the base simulation. The grant did lower the averages a tiny amount ($6 for ages 24-29, $11 for ages 30-35 and $45 for ages 36-41) but this is by less than 0.05 percent in each instance. In the later years, asset levels are higher than the base simulation, because parents are spending less on education, but again, while larger, the increase is not great. For example, by ages 48-53, the average asset level increases by $2,224 or 1 percent. While the pattern of increasing and decreasing is different for the credit experiment, the size of the changes is again very small. Average assets increase by $3 both with and without the cap for households ages 24-29, by $16 and $18 with the cap and without for ages 30-35, and $66 and $69 with and without the cap for ages 36-41. At later ages the assets are lower that the base simulation, as parents spend more on their children's education. However, even by age 48-53, the average assets are only about 0.5 percent lower. The pattern here suggest that while there is a relationship between income, fertility, savings and educational spending, as is seen in the asset accumulation patterns shown in the data and the results section (see figure 6 for example), these particular policies do not have a sizeable impact on the fertility or savings decisions.

The education savings account, on the other hand, does have an influence on fertility and savings, as well as parental contributions and the educational outcomes of children. This experiment shows 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 $1549, up to $7939, while for those without a college degree the increase is $2,233, up to an average of $4,418. This result is most likely driven by the fact that the marginal impact of additional contributions is decreasing. Since more of the lower-educated households are not contributing or saving as much, the marginal impact of their increases in offers is greater, so a lesser marginal incentive to save is needed. Certainly this result is not what is 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 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. This improvement is even greater than that from the grant experiment. However, the cost is also higher. Here the cost is lost revenue from taxable earnings on the education savings. The cost per household comes to almost $8,000. 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 higher contributions, and the associated large cost in lost tax revenue, are due in large part to the fact that the education savings account experiment actually stimulates a large amount of new savings. Unlike the other policy experiments done here, the education savings account greatly alters savings behavior. Average savings for households 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. This clearly indicates that there is additional savings, beyond the amount accounted for just by a higher return, and that it is used by families to help pay for their children's education. The average amounts in the education accounts are also shown in table 10. While they are indeed greater than the increase in savings, especially in later years, they are also not just created from a shift of assets from traditional savings to the higher-return education savings accounts. Along with the increase in savings, fertility rates also rise to where the average number of children 2.410, up from 2.282. Again, this is presumably due to having access to the improved return on savings along with the improved ability to have college-educated children.

7. Conclusion

The model presented here gives structural estimates of 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 model is then estimated with a simulated maximum likelihood procedure using data from the National Longitudinal Survey. The estimated model is then used to run policy experiments to gauge the impact of several recently implemented programs to increase parental support for children's education.

The estimates imply that there is a positive relationship between income, education, having children, and savings behavior. Families with 3 or fewer save more at young ages than those without children, yet married couples without children have greater savings by the time they reach their mid-thirties. The reason is the costs of not only raising children, which are estimated to by quite large, but also of helping to pay for a college education. The model is able to capture the interactions more effectively than if fertility was the result of an exogenous process. The model estimates that the fast 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. Several problems remain with how the model treats the actual education choice of children. First, the education decisions of children are not modeled in a fundamental way. Future work to include this decision more directly, and even 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.) The model here does not directly account for the price of college or the rise in college expenses and their impact. These are just suppresses as part of the estimated outcome process. In addition, there are no difference in the quality of college education here. 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.

Gale, W. G. and J. K. Scholz (1994). "Intergenerational Transfers and the Accumulation of Wealth." Journal of Economic Perspectives **8**(4): 145-60.

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.

Hubbard, R. G., J. Skinner, et al. (1994). "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." Carnegie-Rochester Conference Series on Public Policy **40**(0): 59-125.

Hubbard, R. G., J. Skinner, et al. (1995). "Precautionary Saving and Social Insurance." Journal of Political Economy **103**(2): 360-99.
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.
Lee, J. (2001). Undergraduates Enrolled with Higher Sticker Prices. Washington, D.C., U.S. Department of Education, National Center for Education Statistics.
Leslie, L. (1984). "Changing Patterns in Student Financing of Higher Education." Journal of Higher Education **55**(3): 313-346.
Lino, M. (1998). "Expenditures on Children by Families, 1997." Family Economics and Nutrition Review **11**(3): 25-43.
McFadden, D. (1989). "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration." Econometrica **57**(5): 995-1026.
McPherson, M. and M. Shapiro (1991). Keeping College Affordable. Washington, D.C., The Brookings Institution.
Modigliani, F. and R. Brumberg (1954). Utility Analysis and the Consumption Function: An Interpretation of the Cross-Section Data. Post-Keynesian Economics. K. Kurihara. New Brunswick, N.J., Rutgers U. Press**:** 388-436.
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

SAMPLE SIZE AND AVERAGE NUMBER OF CHILDREN

	Highest Grade Competed		
	High School or less	College	All
No. Obs.	311	245	556
Children (mean)	2.395	2.208	2.313
Children born when 18-23 (mean)	1.225	0.478	0.896
Children born when 24-29 (mean)	0.826	0.971	0.890
Children born when 30-35 (mean)	0.344	0.759	0.527

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

CONTRIBUTIONS TOWARD CHILDREN'S COLLEGE EDUCATION

	No. Obs.	Mean*	Percent Contributing	Mean if Contributing*
All	328	4,273	81.71	5,230
By Parents Education:				
High School or Less	158	2,054	67.72	3,033
College	170	6,336	94.71	6,690

* 1999 dollars

TABLE 4

CHILDREN'S EDUCATION ATTAINMENT

	Percent No College	Percent Some College	Percent 4-year Degree or More
All	42.37	35.73	21.90
By Parents Education:			
High School or Less	52.49	33.72	13.79
College	23.10	39.56	37.34

TABLE 5
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]

λ_{2c}	λ_{2h}	α_{1c}	α_{1h}	α_{2c}	α_{2h}	θ_{1c}	θ_{1h}	θ_{2c}
-0.6286[a]	0.1431[a]	0.0041[a]	0.0053[a]	-0.3699[b]	-0.5472[b]	0.0087[a]	0.0122[a]	-0.5421[b]

θ_{2h}
-0.8246[b]

Income

β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8
11.0244	0.3348	-0.0354	-0.1927	0.1684	-0.0065	0.9626	-0.2011	-0.3001

b_c	b_h
0.9575	0.0088

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]

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

Other Parameters

δ	ψ_c	ψ_n
0.8797	77228	41782

[a]Parameter multiplied by 10^8

[b]Parameter multiplied by 10^{11}

[c]Parameter multiplied by 1000

TABLE 6

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 7

ACTUAL AND PREDICTED SELECTED OUTCOMES

	Actual	Predicted
<u>All Households:</u>		
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
<u>By Education:</u>		
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 7

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 8
POLICY EXPERIMENTS

	Base	Grant	Credit	Credit w/ Limit	Ed. Sav. Account
<u>All Households:</u>					
Parent Contributions to College (mean)	4,364	3,605	5,963	5,610	6,436
Percent of Parents Contributing	70.69	54.51	82.55	82.50	82.56
Childrens Education (percent):					
No College	40.59	35.80	37.99	38.15	34.91
Some College	36.18	37.66	36.65	36.61	37.17
4-year degree or more	23.23	26.54	25.36	25.24	27.92
Children (mean)	2.282	2.285	2.283	2.283	2.410
Total Assets (mean by age):					
24-29	22,292	22,286	22,295	22,295	26,812
30-35	57,925	57,914	57,941	57,943	73,905
36-41	86,193	86,148	86,259	86,262	100,009
42-47	138,729	139,246	138,245	138,306	139,286
48-53	219,224	221,448	218,627	218,942	216,791
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
<u>By Household Education:</u>					
High School Degree or Less:					
Parent Contributions to College (mean)	2,185	1,385	3,866	3,765	4,418
Percent of Parents Contributing	61.87	43.72	80.16	80.06	74.10
Childrens Education (percent):					
No College	52.43	46.10	49.29	49.43	48.03
Some College	34.23	37.13	35.30	35.19	35.91
4-year degree or more	13.34	16.77	15.41	15.38	16.06
Children (mean)	2.370	2.373	2.371	2.371	2.644
Total Assets (mean by age):					
24-29	19,643	19,647	19,647	19,648	29,878
30-35	43,968	43,964	43,990	43,993	57,222
36-41	59,811	59,760	59,921	59,928	76,645
42-47	98,315	98,861	97,553	97,649	102,638
48-53	155,317	156,567	153,905	154,082	144,641
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 8
CONTINUED

	Base	Grant	Credit	Credit w/ Limit	Ed. Sav. Account
Households with College Degree:					
Parent Contributions to College (mean)	6,390	5,817	7,899	7,548	7,939
Percent of Parents Contributing	82.80	69.32	85.84	85.84	90.21
Childrens Education(percent):					
No College	24.00	21.48	22.31	22.46	20.16
Some College	39.00	38.40	38.51	38.59	39.20
4-year degree or more	37.00	40.12	39.18	38.95	40.64
Children (mean)	2.117	2.118	2.117	2.117	2.223
Total Assets (mean by age):					
24-29	25,654	25,656	25,656	25,657	27,011
30-35	75,642	75,624	75,650	75,651	94,099
36-41	119,681	119,644	119,692	119,690	128,734
42-47	190,031	190,511	189,898	189,916	195,500
48-53	300,348	303,807	300,783	301,274	280,819
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

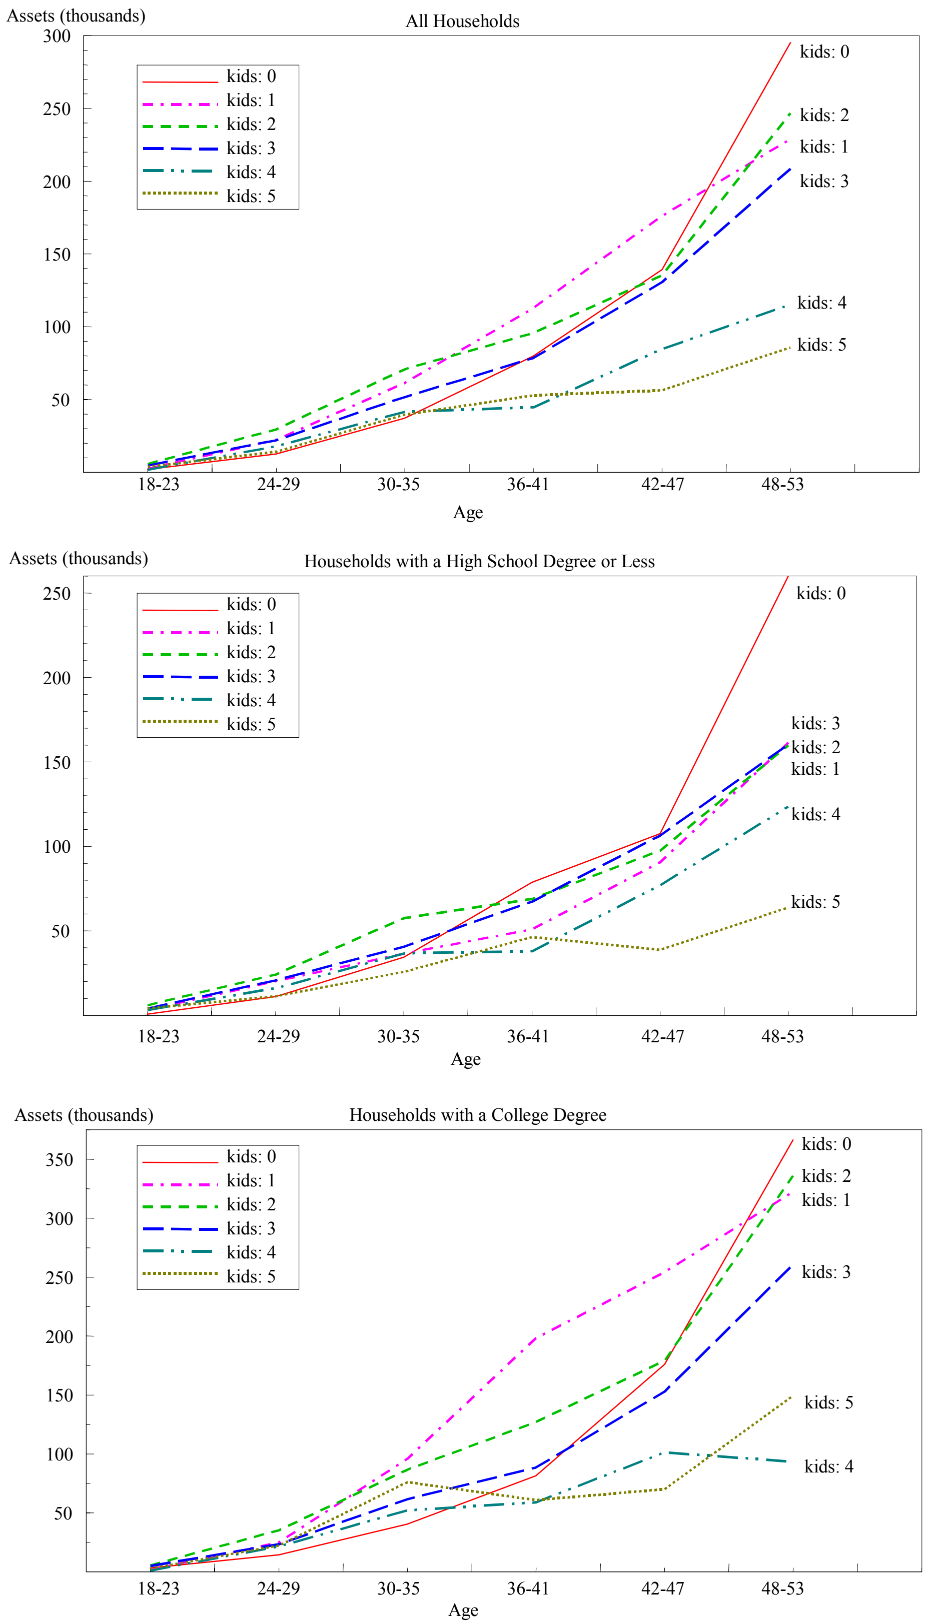

FIGURE 2
ACTUAL AND PREDICTED MEAN ASSET ACCUMULATION

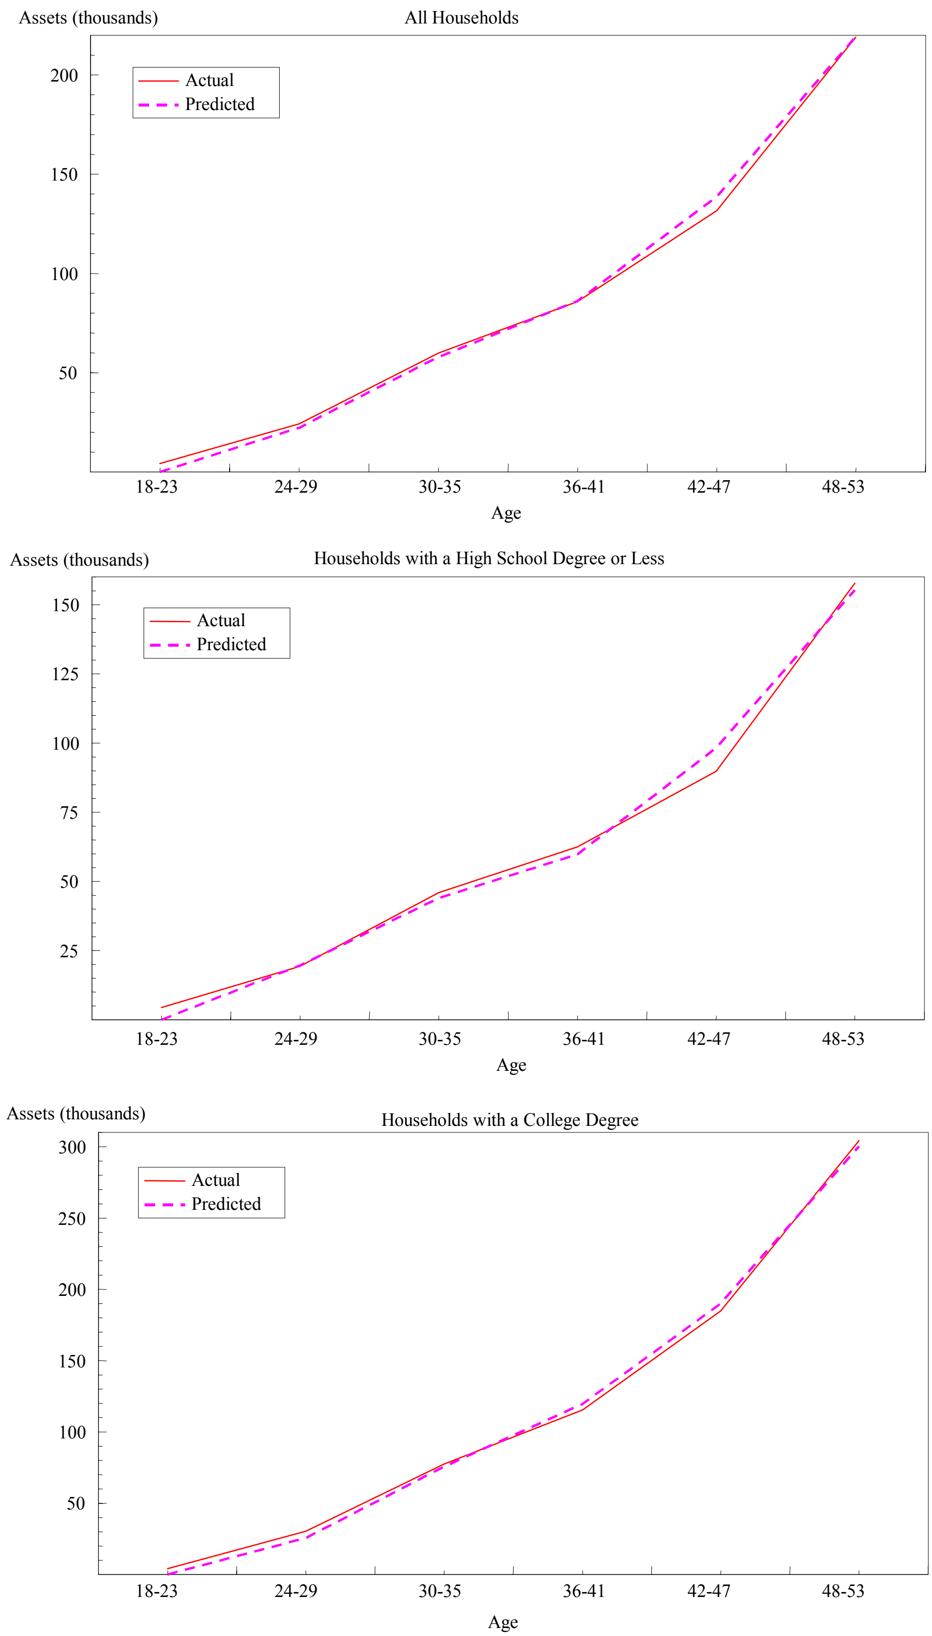

FIGURE 3
SIMULATED ASSET ACCUMULATION BY NUMBER OF CHILDREN

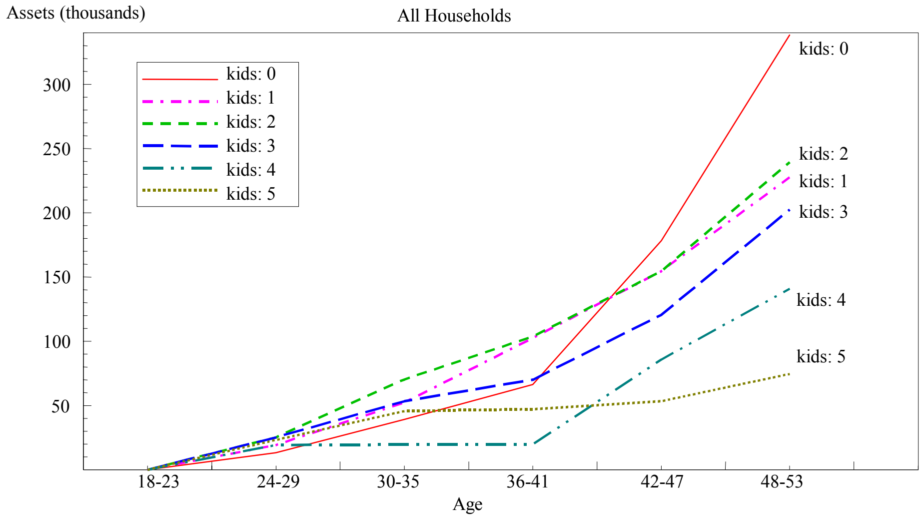

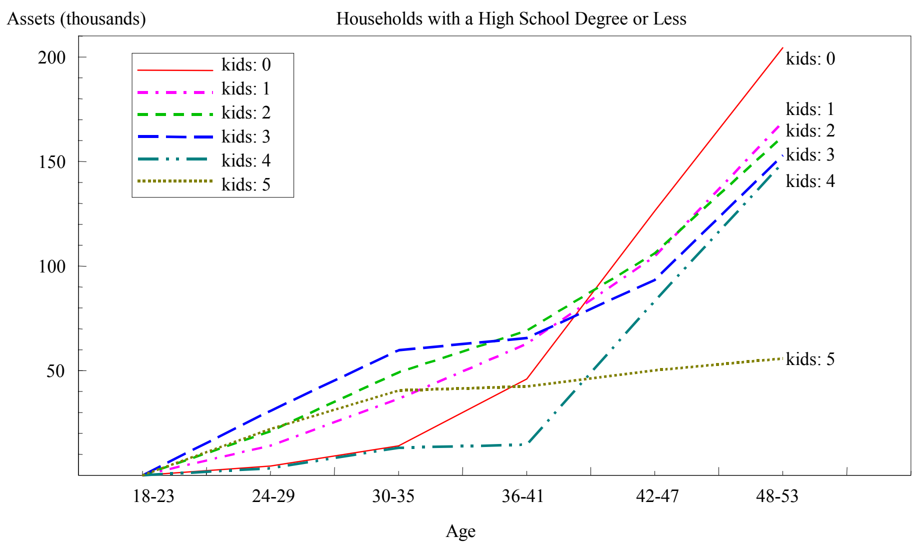

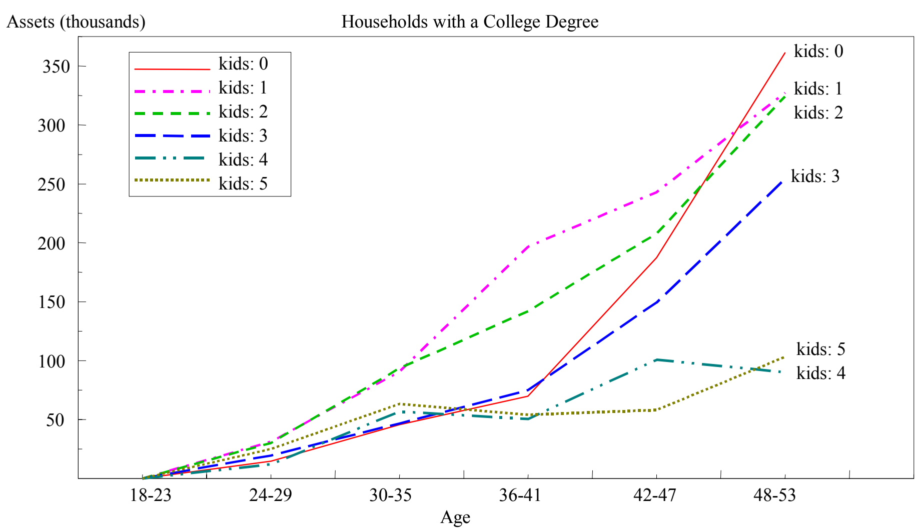

FIGURE 4
ACTUAL AND SIMULATED ASSET ACCUMULATION:
All Households

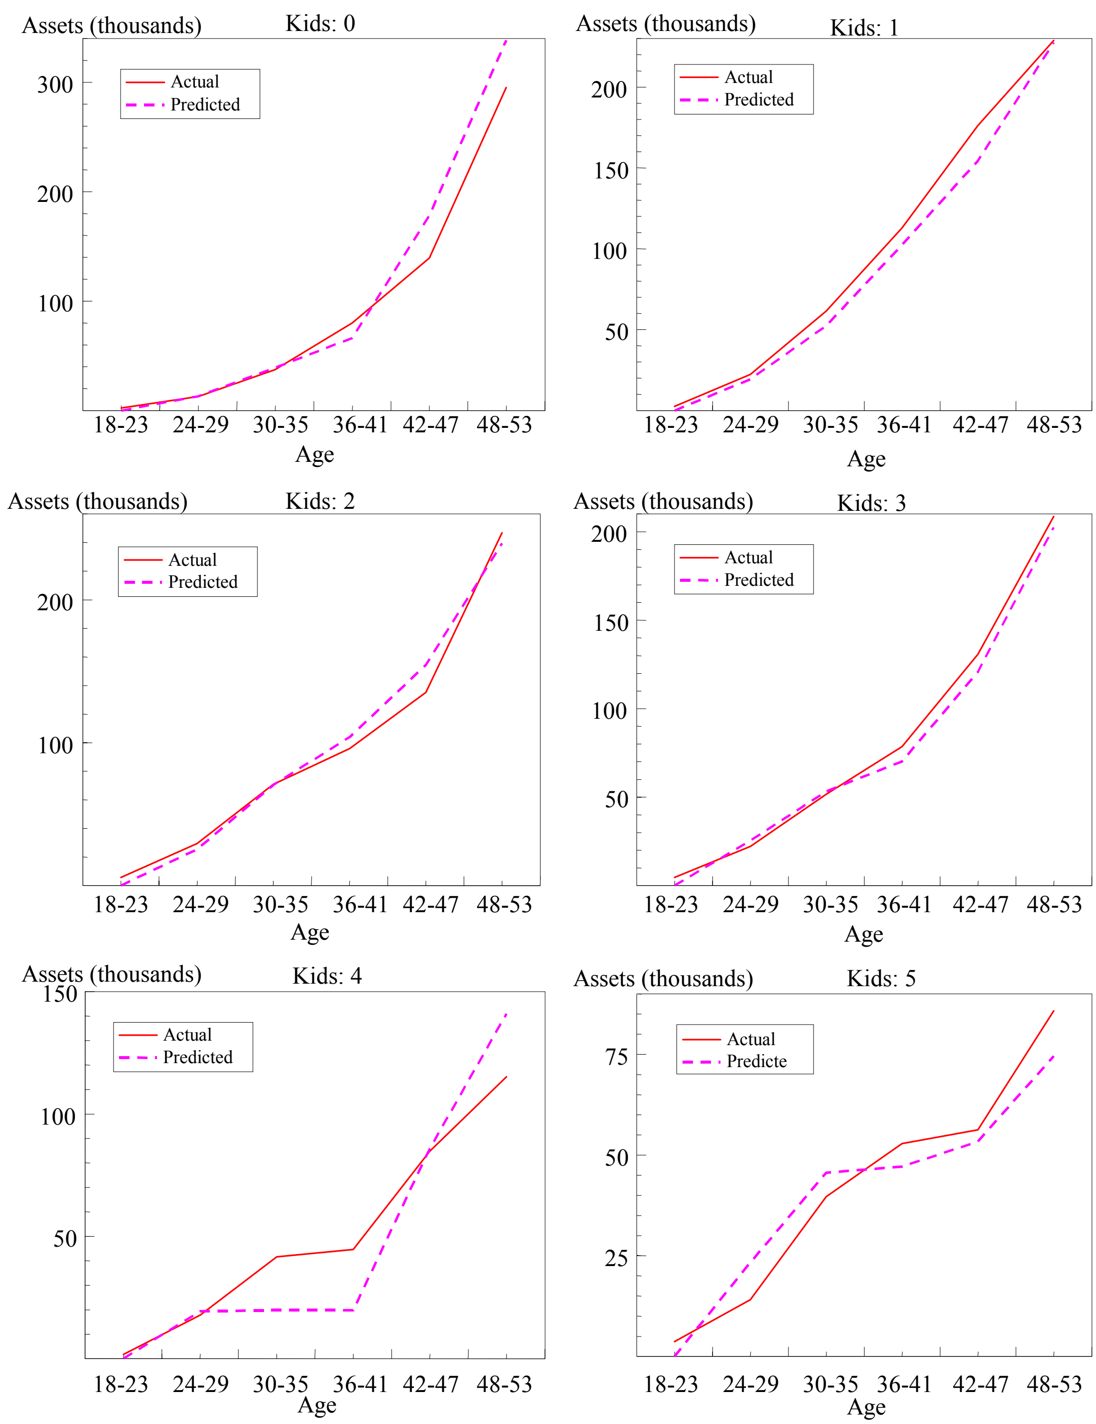

FIGURE 5
ACTUAL AND PREDICTED CUMULATIVE ASSET DISTRIBUTIONS:
All Households