Can Parental Decisions Explain U.S. Wealth Inequality?[*]

John Knowles[†]
Department of Economics,
University of Pennsylvania.

July 1999 (Revised October 2000).

Abstract

Two outstanding puzzles in the literature on the U.S. wealth distribution are the low level of wealth among the poorest half of U.S. households, and the extreme concentration in the richest. This paper asks whether the parental tradeoff betwen quantity of kids and investment per child can resolve these puzzles by generating a similar concentration of wealth as the general-equilibrium outcome in a world where income is partly determined by parental investment in human and physical capital. The paper estimates lifetime inequality from U.S panel data and the model is parametrized to match the empirical relationship between lifetime earnings inequality and family-size decisions. The results suggest that the apparent low wealth share of the poorest U.S. households is due to exclusion of children's human capital from most measures of wealth, and that fertility decisions help explain the concentration of wealth in the richest families.

[*]Based on Chapter 2 of my PhD dissertation, prepared for submission to the Economics Dept. of the University of Rochester, under the supervision of Per Krusell. Thanks to Hanan Jacoby and Sylvain Leduc for helpful discussions, and to Per Krusell for continuing guidance. This version is very preliminary, so comments are particularly welcome, but please check with me before citing my results.

[†]Email: **jknowles@econ.sas.upenn.edu.**

1 Introduction

Two dimensions of wealth inequality that are difficult for current theory to explain are the lack of wealth accumulation among the poorest 40% of households, and the extreme concentration of wealth among the richest 1% of households, as noted by Wolff (1994) and Diaz-Gimenez, Quadrini, and Rios-Rull (1997). Explanations based on the precautionary and life-cycle savings motives were rejected by Huggett (1994), who showed that an equilibrium model of life-cycle savings with various sources of uncertainty could not replicate these dimensions of wealth concentration. This failing was also noted by Aiyagari (1994) in an equilibrium analysis of the precautionary savings motive.

More recently, Krusell and Smith (1996) proposed an explanation of equilibrium wealth concentration based on the assumption that wealthier agents are more patient. Castaneda, Diaz-Gimenez, and Rios-Rull (1998) on the other hand get higher wealth concentration by assuming that bequests are the only way to invest in one's children. In these papers, earnings inequality is fixed exogenously and savings are the only investment margin; this paper by contrast proposes an explanation of wealth concentration based on the idea that wealth and earnings inequality are jointly influenced by the interaction among several parental decisions.

Parental decisions in real life are made subject to at least two forms of market incompleteness: uncertainty about the child's eventual wage is uninsurable, and parents can not borrow against their children's future earnings. It is well known that these missing markets, combined with the assumption that children are more costly for higher-income parents, can generate the negative relationship between parental income and fertility, and the positive intergenerational correlation of wealth and earnings. The implied negative relation between fertility and children's outcomes is the result of several trade-offs faced by parents: between their consumption and their altruistic concern for their children, between quantity and "quality" (welfare) of children, and between children's human capital and financial bequests.

These trade-offs may therefore result in a funneling effect of fertility on wealth: rich parents invest more income in fewer children, and this investment is more biased towards non-human capital than that of poor parents. Another reason why differences in fertility over the income distribution may play a key role, given that income differences are correlated across generations, is that the weight of the descendants of a given class of parents in future generations is proportional to the parent's fertility. Hence parental decisions may be part of a candidate explanation along both dimensions of the wealth-concentration puzzle; investments in children's human capital are likely to have a higher return than financial assets for poor parents, thus explaining their low holdings of other assets, while the fertility-earnings relationship may explain the high level of accumulation among the richest households[1].

This paper uses these elements to develop a dynamic equilibrium theory of income inequality and compares the quantitative implications of the theory for the wealth distribution to measures of lifetime inequality in the U.S. The theory consist of two components: a model of parental decisions regarding fertility and investment in children, taking prices as given, and a general-equilibrium environment where these parental decisions determine the capital-labor ratio, and hence the interest rate and the real wage. While the decisions analysed here have been modeled in the previous literature, there has been little work examining these margins jointly, particularly in an equilibrium context with heterogeneous parents. The argument of this paper is that it is the interaction among these margins that is crucial for understanding long-run wealth inequality in the U.S.

The equilibrium theory of the income distribution consists of the stochastic steady state of an economy populated by a continuum of such parents who differ along two dimensions: their realized wage, which depends on their luck in the labor market, and on the bequest they inherit from their parents. Because this setting allows for heterogeneity along several important margins that are also observable in cross-sectional data, such as fertility, wealth and income, the model is amenable to quantitative analysis of equilibrium inequality, in the spirit of recent analyses of the U.S. wealth distribution by Aiyagari (1994) and Huggett (1994). The equilibrium analysis is an essential part of the story; for while the decision model may be plausible, it is not clear how consistent the different components may be with each other in an equilbrium wealth distribution. For instance, since high-income people have more wealth, and children are normal goods, then the theory may generate a positive fertility-earnings relation, very low intergenerational correlations of earnings, or even a negative relation between earnings and wealth.

The goal of the empirical analysis is to construct analogs of the inequal-

[1] This motive for savings has recently been explored by Castaneda, Diaz-Gimenez, and Rios-Rull (1998). Their model, which combines bequest and precautionary savings motives, can indeed match the U.S. wealth distribution. However the model takes the earnings distribution as given, and abstracts from the other parental decisions considered here.

ity dimensions in the model. Regression-based estimates are used to predict lifetime household income for a representative sample of women in the Panel Study of Income Dynamics. On the basis of these lifetime income measures, this section reports inequality statistics and the relationships between lifetime income and outcomes such as fertility, reliance on transfers, and children's lifetime income. The empirical results show that the Gini coefficients for earnings and total income are much lower than is the case for annual measures, the inter-generational earnings correlation is close to that estimated for fathers and sons in the previous literature, and that fertility is declining in household earnings. A benchmark version of the model is constructed by specifying the model's parameters so as to obtain a close fit between the model's steady-state distribution and those features of the data pertaining to lifetime earnings inequality and fertility decisions.

Overall, the principal results of the quantitative analysis are that bequest wealth is highly concentrated, that endogenous fertility is responsible for most of this concentration, and that the level of savings generated by bequests is quite low when fertility is endogenous. The stationary equilibrium implies significantly more concentration of wealth than in the models that rely on other savings motives; most of the population hold all their wealth in the form of children's human capital. The funneling effect of the fertility-income gradient is not sufficiently strong to replicate the concentration of wealth in the richest 1% of the population, who hold 22% of total wealth in the model, but the model does remarkably better in this dimension than papers based on lifecycle or precautionary savings motives. The distributional statistics that the model is capable of matching include the intergenerational earnings correlation, the relation between average fertility and earnings, and the Gini coefficients for earnings and wealth. The savings rate is 3%, roughly equal to that attributed to precautionary motives by Aiyagari (1994).

Suppressing fertility variation results in a much higher average savings rate, and in much less concentration of wealth among the very richest households; the holdings of the richest 1% declines to 7% of total wealth. Furthermore the level of savings more than doubles, even though earnings uncertainty is held constant. These results suggest that family motives are an important explanation of the wealth concentration in the U.S., and should be studied in the context of other wealth-accumulation motives that have drawn more attention in the literature on wealth inequality in the U.S.[2]

[2]Recently Veloso (2000) has undertaken a similar excercise for the case of Brasil.

The decision theory in this paper is based on previous models, but is among the first to combine all of these margins in a stochastic environment. Thus the quality-quantity tradeoff is due to Becker and Tomes (1976), while the analysis of the allocation of parental investment between human and physical capital draws on Becker and Tomes (1986). The dynastic decision framework is based on Becker and Barro (1988) and Alvarez (1994); the latter integrates these elements in a stochastic environment. The equilibrium modeling of the income distribution is in the spirit of Loury (1981), who analyses an economy with human capital investment, and Laitner (1992), who considers the case of savings with exogenous human capital. Finally the empirical analysis refines the work of Solon (1992), Zimmerman (1992) and Mulligan (1993) by susbstituting OLS projections of income for the simple averages used in the previous literature.

The next section presents a formal model of the income distribution, while Section 3 analyzes some general properties of the model. Section 4 discusses how the model is calibrated and solved. Results are discussed in section 5, and the findings of the paper are summarized in Section 6.

2 A Model of Income Inequality

2.1 The Environment

Consider an overlapping-generations economy populated by a continuum of two-period-lived agents who are identical *ex ante* and reproduce asexually. Agents live the first period of their lives as children, when they make no economic decisions, and the second period as parents or adults. In the second period of their lives, agents work and raise n children, deriving utility from their own consumption c and also from the expected welfare of their children as adults. At the beginning of the second period of life, adult agents experience market-luck shocks z that interact with their education level e to determine their level of labor skill h.

The preferences of agents who are adults at time t are represented by a utility function U_t. Parents care about their own direct utility $u(c)$ from consumption and about the welfare $\{U_{t+1,i}\}_{i=1}^{n}$ of each of their n children. Activities of children in the first half of life do not enter preferences. As in Becker and Barro (1988), it is assumed that parents weight welfare per child

$U_{t+1,i}$ by a decreasing, concave function $b(n)$ of family size[3]: thus the welfare of a parent at time $t = 0$ can be written recursively as:

$$U_0 = u(c) + b(n) \sum_{i=1}^{n} E\{U_{1i}|z\} = u(c) + b(n) nE\{U_1|z\} \qquad (1)$$

The key assumptions about these functions are that $u(c)$ is twice continuously, differentiable, increasing and concave, and that $b(n)$ is decreasing in n. The derivative $u'(c)$ satisfies Inada conditions: $\lim_{c\to 0} u'(c) = \infty$ and $\lim_{c\to\infty} u'(c) = 0$.

These assumptions imply that parents have children in order to smooth consumption over more agents: despite their altruism, parents would not have children if they did not experience diminishing marginal utility of consumption. The greater the curvature of $u(c)$, the larger the family. The form of the discount factor function $b(n)$ ensures that parents do not all choose the maximum feasible family size: the greater the number of children, the less the welfare of each child contributes to parent's welfare.

Parental income $y(e, a, z)$ is the sum of labor income and capital income from inheritances. It depends on the parental education e, their inheritance a, and their productivity shock z.

Raising children requires inputs of both goods and parental time. The cost of children includes an income-dependent component that is independent of the investment per child and is increasing in parental income. For simplicity, it is assumed that the cost ϕ of producing a child of a given education level is independent of the family size and that the time input is a constant θ per child. Let the function $\phi(e', e, z)$ represent the cost of children's education e'; the cost function $\phi(e', e, z)$ is increasing in children's education e', decreasing in e, and the derivative $\phi_1()$ satisfies $\lim_{e'\to 0} \phi_1() = 0$ and $\lim_{e'\to\infty} \phi_1() = \infty$. These assumptions ensure that all children will receive some human capital, and that high levels of education are extremely costly. The time cost θ of children represents the portion of child-raising time that can not be substituted for by market services. The effect of θ on the model is to make the child-raising cost increasing in the *ex post* wage of the parent.

Heterogeneity across families is reinforced in this model by stochastic variations across generations in the market luck z of the children, which follows a first-order Markovian process and is independently distributed across

[3]These assumptions on the altruism function ensure diminishing marginal returns to number of children.

families. Children of the same family share the same realization. The probability that a child whose parent is in state z enters state z' is given by $\pi(z, z')$ =prob$(z_{t+1} = z' | z_t = z)$. Together, market luck and education determine the adult's skill level $h(e, z)$. The key assumption here is that $h_1() > 0$, so that increasing e is assumed to shift the distribution of h to the right in the usual sense of first-order stochastic dominance.

Parents are endowed with one unit of time. The time allocation of a parent who works l units of time and has n children at education level e' must satisfy the following constraint:

$$1 \geq l + n\theta \tag{2}$$

as well as a feasibility constraint on fertility: $n\theta \leq 1$.

Output is produced via a neoclassical production function $F(K, L)$ that gives output as a function of the aggregate labor input L and aggregate capital services K. The total goods available for consumption in the economy is therefore given by $Y = F(K, L)$. The production function satisfies constant returns to scale and the marginal products F_x satisfy $\lim_{x \to 0} F_x() = \infty$ and $\lim_{x \to \infty} F_x() = 0$ for $x \in \{K, L\}$. The current period's capital stock is equal to the aggregate savings of the previous period: capital depreciates completely over a generation.

2.2 The Market Arrangement

The key assumption about the markets in this economy is that there are two forms of incompleteness: parents can not borrow against the future income of their descendants, and they can not insure their children against bad luck in the labor market. There is a risk-free asset; the parent's holdings of this asset are denoted by a. The restriction that parents can not borrow against income of future generations implies that $a' \geq 0$. After learning their own labor-luck realizations, parents make all their decisions. They choose fertility n, the education level of their children e', bequests to children a' and consumption c. From the time constraint (2) ,the fertility choice implies the agent's labor supply $l = h(e, z)(1 - n\theta)$.

The wage w per unit of labor supply and the interest rate r on the risk-free asset are given by the marginal products in the production function: $w = F_L(K, L)$,and $r = F_K(K, L)$.

Transfers from social programs that redistribute income from rich to poor families are an important component of income for poor families in the U.S, and are represented in the model by a function $\delta(y)$ that gives the level of transfer a family receives as a function of it's pre-transfer income y. This transfer program is paid for each period with a proportional income tax at rate τ.

The aggregate state of the economy is given by the density function $\mu(e,a,z)$, which gives the proportion of adults in state (e,a,z). The total mass of the population doesn't matter for prices. Given that there is no aggregate uncertainty, and given parental decision rules $g^x(e,a,z,\mu)$ for each choice variable $x \in \{n, e', a', c\}$, parents can infer perfectly the future paths of prices and aggregate quantities from the density function μ.

Let the income of a parent in state (e,a,z) be written as:

$$y(e,a,z,\mu) = w(\mu)\, h(z,e)\,(1 - g^n(e,a,z,\mu)\,\theta) + r(\mu)\, a \tag{3}$$

. Then the Bellman equation for the parent's problem is given by:

$$V(e,a,z,\mu) = \max_{c,n,e',a'} \{u(c) + b(n)\, nE\,[V(e',a',z',\mu')\,|z,\mu]\} \tag{4}$$

subject to the budget constraint:

$$\delta(e,a,z,\mu) + y(e,a,z,\mu) - c - n\,[\phi(e',e,z,\theta) + a'] = 0 \tag{5}$$

and the usual non-negativity conditions: $c \geq 0, n' \geq 0, e' \geq 0, a' \geq 0$.

The Bellman equation says that, when the distribution of all parents in the economy is given by μ, the value to parents of being in state (e,a,z) is equal to the utility from the optimal choice of consumption plus the expected value of the n children given the optimal choice (e',a') of education and bequest, discounted by the factor $b(n)$. The expected value of the children depends on the parent's realization z because this influences the distribution of their idiosyncratic productivity shocks z', and also on the aggregate state μ, which determines the wage and interest rate next period.

The non-standard aspects of the dynamic programming problem are that the discount factor is endogenous and that some of the choice variables enter multiplicatively into the budget constraint. These features mean that standard methods cannot be used to demonstrate concavity of the value function and the existence of optimal policies[4].

[4] Alvarez (1994) considers a simpler version of this problem, in which the cost of children

2.3 The Recursive Competitive Equilibrium

Since parents save only for bequest purposes and since capital is equal to aggregate savings of the previous period, then current capital is equal to the sum of bequests received by the current generation of adults. The level of aggregate capital services is therefore given by:

$$K(\mu) = \int a d\mu(e, a, z) \tag{6}$$

Labor supply per capita, on the other hand, depends on the fertility-decision rules of the parents:

$$L(\mu) = \int h(e, z) \left[1 - g^n(e, a, z, \mu) \theta\right] d\mu(e, a, z) \tag{7}$$

Aggregate demand is the sum over all agents of desired consumption and spending on children:

$$Y^D(\mu) = \int \left[g^c(e, a, z, \mu) + g^n(e, a, z, \mu) \phi\left(g^e(e, a, z, \mu)\right)\right] d\mu(e, a, z) \tag{8}$$

Definition 1 *A* ***recursive competitive equilibrium*** *of this economy consists of the following objects: a pair of price functions* $\{w(\mu), r(\mu)\}$*, a list of decision rules for fertility* $g^f(e, a, z, \mu)$*, education* $g^e(e, a, z, \mu)$*, investments of child-raising time* $g^\theta(e, a, z, \mu)$ *and goods* $g^\phi(e, a, z, \mu)$*, financial bequests* $g^a(e, a, z, \mu)$*, labor* $g^\tau(e, a, z, \mu)$*, and consumption* $g^c(e, a, z, \mu)$*, a per-capita goods-demand function* $Y^D(\mu)$*, a pair of factor demand functions* $\{K^D(\mu), L^D(\mu)\}$ *representing demand for capital and labor, respectively, and a law of motion* T *for the measure* μ*, such that the following conditions are satisfied for every possible aggregate state:*

---

is not dependent on the parental state. He shows that this type of problem can be rewritten as the problem of an "extended family" in which all descendants of a common ancestor pool their income. The non-standard features of the original problem disappear in the extended-family problem, the solution of which requires a relatively minor extension of standard techniques to the case of a homogenous return function with constant returns. The solution of the extended-family problem implies that the value function that solves the original problem (4) is unique and the optimal policy correspondence non-empty. If in addition, the value function is concave, then the consumption policy is continuous in the state variables, and therefore so is total spending on children. Furthermore, concavity of the value function also implies an upper bound on capital accumulation (Aiyagari (1994), Proposition 4). Continuity of the education-investment and bequest policies, however, is not guaranteed because of the interaction with fertility in the budget constraint.

1. Optimality:

   (a) The parental policy solves the parent's dynamic programming problem (4) given the prices (w, r).

   (b) The prices (w, r) equal the marginal products of labor and capital, respectively.

2. Feasibility:

   (a) The goods market clears: $Y^D = F(K(\mu), L(\mu)) + \delta$, where Y^D satisfies (8).

   (b) The factor markets clear, given parental polices: $K(\mu)$ and $L(\mu)$ satisfy conditions (6) and (7), respectively.

3. Budget balance: the cost of transfers equals the revenue collected from taxes:

$$\int [\delta(e, a, z, \mu) - \tau y(e, a, z, \mu)]\, d\mu(e, a, z, \mu) = 0.$$

4. Consistency of the the law of motion T of the probability measure with individual decisions: for any $S_1 = E_1 \times Z_1 \subseteq S$ and set $A_1 \subset \mathbb{R}_+$:

$$T\mu(S_1, A_1) = \frac{1}{M'} \int_{z' \in Z_1} \int_{G(S_1, A_1, z)} g^f(e, a, z, \mu)\, d\mu(e, a, z, \mu)\, \pi(z, dz') \tag{9}$$

where

$$G(S_1, A_1, z) = \{(e, a, z, \mu) \,|g^e(e, a, z, \mu) \in E_1\} \cap \{(e, a, z, \mu) \,|g^a(e, a, z, \mu) \in A_1\}$$

and M' is the mass of the current children's generation relative to that of the adults:

$$M' = \int g^f(e, a, z, \mu)\, d\mu(e, a, z, \mu)$$

These conditions are quite standard: Equation (9) says that the measure next period of parents in a given region of the state space is equal to the mass this period of parents who choose for their children education and transfers in that region, weighted by the fertility of the parents and by the probability that the children's luck realization is also in the same region. Dividing through by the average fertility rate M' ensures that $\mu(e, a, z)$ is a probability measure; the mass of the population is irrelevant in this environment.

Definition 2 *A* ***stationary (steady-state) equilibrium*** *is a recursive competitive equilibrium where the distribution of agents is constant over time. This implies that the wage and interest rates are also constant, as are the capital stock per-capita, and the growth rate of the population.*

The rest of the paper restricts attention to stationary equilibria, so the aggregate state is suppressed when referring to the value function and decision rules of the agents. The parent's dynamic problem can then be written as:

$$V(e, a, z) = \max_{c,n,e',a'} \{u(c) + b(n) nE[V(e', a', z') | z]\} \tag{10}$$

subject to the constraints 5 and 2.

2.4 Parental Behavior in the Model

Although the principal motivation for parents in the model is altruism for their children, the mathematical representation of this motive is that parents have children in order to spread consumption across more utility functions. The higher the curvature of the utility function, the more children parents will tend to have. Because having children is costly, however, wealthier parents will tend to have more children than poor parents, holding constant the cost of children, because the benefit of an additional child is greater the higher is consumption per person. As the wealth of parents increases, holding earnings constant, fertility will also increase.

This suggests that in order to generate a fertility-earnings relationship that is consistent with the observation that high-income parents tend to have fewer children, it is necessary that the costs of children are higher for those with high earnings. This effect is generated via two channels in the model. Children impose significant time costs on parents,which implies that the opportunity costs of children is higher for high-wage parents, and the pecuniary cost of children is increasing in parental income. Both assumptions have been widely used in the literature.

Parents with low wages relative to their goods endowment will not be affected much by this time cost, however; their fertility is more likely to be limited by the fixed portion of the goods cost of raising children. Neither of these two categories of fertility costs are likely to affect parents with large non-labor incomes; in general the effect of non-labor income is to increase fertility. One type of fertility cost that all parents face, regardless of income, is the loss of utility per child that accompanies an increase in the number of children. This is due to the assumption of concavity of the discount factor function.

Given the homothetic structure of parental utility, the share of income that parents allocate to investment in children is roughly constant, as explained in Alvarez (1994). The departures from this rule are due to the fixed costs of children in the case where the choice of children is discrete.

The implications of the model for the allocation of investment are best understood by assuming that a parent has chosen a fertility level n and consumption c. The remaining decision is how to allocate the residual income over investment in children's human capital and savings for bequests. If the solution involves positive amounts of each, then the following condition must be true at the optimum:

$$E\left[V_e\left(e', a', z'\right)|z\right] = E\left[V_a\left(e', a', z'\right)|z\right]\frac{\partial \phi\left(e'; e\right)}{\partial e'} \tag{11}$$

This condition says that the marginal benefit of an extra unit of spending on human capital equals the marginal benefit of an additional unit of bequest per child. The expectations in equation (11), according to the envelope theorem are given by the following equations:

$$\begin{aligned} E\left[V_e\left(e', a', z'\right)|z\right] &= E\left[u'\left(c'\right)\left(1 - n'\theta\right) z'|z\right] && (12)\\ E\left[V_a\left(e', a', z'\right)|z\right] &= E\left[u'\left(c'\right)|z\right] R && (13)\end{aligned}$$

From Equations (11) and (12), it is clear that the expected fertility of the children lowers the marginal benefit of education. Since bequests increase the expected fertility of the children, it is clear that human capital is not strictly increasing in wealth: once parents start making bequests, they will reduce the human-capital investment in their children. It is also clear, since the marginal cost at zero human capital is zero, that if the expected labor supply of the children is bounded below away from zero, and the labor luck process bounded below so that strictly positive human capital implies a strictly positive wage, then the rate of return on human capital investment is infinite at zero, so all parents will invest in human capital. The rest of this paper assumes that fertility is bounded above by the same maximum number of children $\bar{n}$ for all households. A sufficient condition for the first requirement is that this upper bound on fertility $\bar{n}$ be small enough so that the total time cost of raising $\bar{n}$ children is less than the time endowment per parent.

Only the richer parents will save for bequests, however, because the borrowing constraint implies that for the poorest parents, the rate of return on human capital is higher than the interest rate. This insight, due to Becker and Tomes (1986), is illustrated in Figure 1. At wealth level a_0, parents who spend less than I_0 per child will invest only in human capital because the percieved rate of return, $E[V_h/V_a]$, is higher than that on physical capital, R. Furthermore, if parents do make bequests to children, say $a_1 > 0$, then they will invest a lower amount in their children's human capital, $I_1 < I_0$. This is partly because bequests raise expected fertility, hence reduccing the time the children are expected to spend in the labor market. The children with the most education will be those whose parents were just short of the level of wealth per child at which parents start making bequests.[5]

Equations (11) and (13) combine to give the following necessary condition for bequests:

$$\frac{u'(c)\,n}{E\left[u'(c')\right]} = b(n)\,R$$

As the inheritance of agents increases, labor income, and hence uncertainty, become less important for two reasons: first, because fertility approaches the upper bound, and hence labor time is reduced, and second, because as capital income grows, the maximum attainable labor income becomes an increasingly small fraction of total income. This implies that the parental problem approaches asymptotically the simple savings problem with one asset with a non-stochastic return. For homothetic preferences, the solution to the asymptotic problem implies that bequests are a constant proportion of

[5] One implication of this interaction between human capital and bequests is that the policy rules for investment are likely to be discontinuous. If fertility is limited however to a selection from a finite set, then this discontinuities occur at a finite number of points.

the inheritance[6]. Therefore as parents become richer, their bequests will approach this proportion from above, since they will also make bequests from their labor income and consumption endowment.

For a steady-state equilibrium to exist, the bequests must be bounded above. A sufficient condition for this in the asymptotic problem is the following upper bound on the interest rate:

$$\frac{b(\overline{n})}{\overline{n}} \geq R$$

When this condition is satisfied, there exists an inheritance level $\overline{a}$ such that for any inheritance $a > \overline{a}$, the optimal bequest will be less than $\overline{a}$.

3 Quantitative Analysis of U.S. Income Inequality

It is well known that the various dimensions of the model outlined in the previous section can provide simple explanations of the analogous qualitative features of inequality or intergenerational mobility, at least in simple models that ignore the other dimensions of inequality. The question is whether it is possible for the model outline in the previous section to simultaneously match the basic *quantitative* features of U.S. income inequality that correspond to features of the model's equilibrium steady state, in particular the joint distribution of fertility and earnings. To answer this question, this section proceeds by estimating empirical features of life-time inequality and fertility from U.S. panel data, and then specifying the model so its steady-state distribution matches those features. The section concludes by analyzing the implications of this specification for the other features of inequality, principally non-labor wealth.

In the next subsection, some empirical features of U.S. inequality are derived from analysis of the Panel Study of Income Dynamics. The following sub-section gives the functional forms and parameter values chosen for the benchmark model; this is followed by a description of the computational algorithm, and finally by a comparison of the model's results to empirical features of U.S. inequality.

[6] For CES utility, the asymptotic problem can be solved analytically for the bequest as a proportion of the inheritance.

3.1 Lifetime Inequality in the U.S.

The goal of this subsection is to determine the features of U.S. income inequality that the model developed above should be able to explain. The idea is to calculate empirical analogs of the model's output from the Panel Study of Income Dynamics, which contains 25 years of household income and demographic information for a representative sample of the U.S. population. These analogs include measures of the shape of the earnings distribution, such as the coefficient of variation and the Gini coefficient, as well as the relationship of average fertility to earnings and the intergenerational correlation of income.

Conceptually, the challenge is to map the real-life households found in the data on to the relatively abstract families defined by the model. Two important dimensions of the data which are missing in the model are gender and the adult life-cycle. Gender matters here in several ways, but these can be reduced to the question of how a household is defined; this study will define households as consisting of adult women who are household heads or spouses. The advantage of focusing on women is that they are empirically more likely to remain with the children in case of household instability, and, equally important, the data on women's fertility in the PSID is more complete than is the case for men's.

Income is defined as the sum of the income of the household head and the spouse. Since female-headed households tend to be very poor, while single-males tend to be relatively well-off, this decision implies that income inequality will be higher than would be the case if the sample were restricted to male-headed households. Another implication of the model's structure is that fertility predicted for a class of households by the model should equal half the number of children observed in the data, since real-life children each have two parents.

The adult life-cycle in the model consists of one period; in the data, if we think of adults as those people aged 26 or more, then we are left with around 40 periods of labor income, often more. To map the data onto the model therefore requires compressing the data. The way that is done here is to consider earnings to be the present value of future labor income, discounted back to age 30. An important technical problem associated with this strategy for mapping panel data onto the model is that the PSID has at most 25 years of data (1967-1991) for any one household, so realised lifetime income is not observable. Recent studies of lifetime income such as Solon (1992) and Mulligan (1993) tend to use averages of annual income as a measure of lifetime income, but Knowles (1999) shows that this method understates the persistence of income across generations.[7] Lifetime income will be estimated therefore using the same technique as in Knowles (1999), applied to females rather than males.

This method involves predicting annual household income, conditional on the characteristics of each female, by using the estimates derived from an income-dynamics model with individual fixed effects. Because we can not observe the marital status of the women after the data set ends, we can not condition on the husband's characteristics, or any other characteristics of the women that are not observable after the age of the youngest sample members at the time of the most recent observation in the data set. Since males tend to contribute a greater share of income to households, this implies that, unless there is perfect marital assortment by earnings potential, the observable characteristics of the female will not be as good a predictor of household earnings as those of the male, so the results are likely to be less precise than those of Knowles (1999). An important advantage however is that this method removes the inter-cohort and cyclical variation in labor income, allowing comparability over women from different cohorts.

The basic assumption is that the log of annual household income y_{it} follows a deterministic pattern with annual disturbances ε_{it} drawn from a log-normal distribution. Income depends on a vector x_{it} of polynomial functions of the individual's potential work experience. The specification also allows for individual fixed effects α_i and for time-varying effects δ_t of individual characteristics w_i, where D_t represents year dummy variables :

$$y_{it} = \alpha_i + (\beta + d_i)\, x_{it} + \delta_t D_t w_i + \varepsilon_{it} \tag{14}$$

. The characteristics w_i are assumed constant (e.g. education, age cohort and race); those that vary deterministically (eg age, age of spouse or children) are included in x_{it}, while those that vary stochastically with time are not inferrable outside of the dataset and hence cannot be used to predict income in the future. The non-time-varying effects of an individual's observable

[7] This is because the standard method implicitly assumes that all individuals have the same age-income profile. It is well-known that people with higher education tend to have steeper profiles than the less educated, which implies that the future income of the young college-educated workers tends to be underestimated.

characteristics are accounted for by the individual fixed effect, which will also reflect the effects of unobservable heterogeneity among individuals. The process for the error ε_{it} is assumed to be independent of time and identical for all individuals.

Lifetime income is calculated as the sum of the estimated annual income purged of the cohort and year effects, and discounted annually by the factor 0.95. This involves removing cohort effects from the individual fixed effects, and predicting annual income on the basis of this residual fixed effect and the estimated regression coefficients for effects that do not depend on time. As a result, the predictions are conditional on age and interactions of other variables with age, but not on interactions with time. Because time and experience are co-linear for individuals, the time variable is replaced with aggregate indicators, in this case the U.S. unemployment rate.

To estimate the distribution of income by this procedure therefore requires two steps: estimation of the income dynamics, and construction of lifetime income. The data sample for the first step consists of annual observations on all women from the representative cross-section portion of the PSID who were between ages 10 and 80 in 1968 and who reported the number of children they had. Only observations made while the women were between 30 and 80 years old were used. This resulted in a sample of 37, 411 observations drawn from 2183 women. The characteristics of this sample are given in Table A1 in the appendix. The average age in 1968 was 33. Nearly all (98%) of the sample completed primary school, 80% completed high school and 31% completed college. [8]Women from households where the head in 1968 was black make up 9 percent of the sample. The average number of children born per woman is 2.50. Total income per household, in 1967 dollars, averages $10,081, while the average labor income is $8,372.

The dependent variables in the income dynamics regression were total household income, and household labor income. Total income is defined as the sum of money income of the household and head and the spouse and money saved through food coupons, deflated by the CPI. The explanatory variables include household characteristics, such as race and education, interacted with age and with aggregate variables, in order to distinguish the effects of business cycles and secular variation from the age-income profile.

The income model (14) was estimated on the income dynamics sample by

[8]The PSID reports years of education, not attainment. This variable was converted to levels using the standard length of each phase of education.

OLS, using dummy variables for each person to reflect the individual fixed effects. The results are reported in the appendix, Table A2, for the log of annual labor income and the log of annual total income.

Table 1: Lifetime Income Distribution*

Statistic	Women's Income	
	Total	Labor
Mean	$224,541	$170,136
Std. Dev.	$152,394	$134,457
Gini Coefficient	0.32	0.38
Percentile/Mean		
1	0.17	0.06
5	0.28	0.18
10	0.38	0.30
25	0.57	0.51
50	0.89	0.85
75	1.26	1.27
95	1.99	2.17
99	3.31	3.97

*In 1967 $, as estimated by author from cohort aged 10-30 in 1968 in PSID 1967-1991.

The estimated lifetime income distribution was based on the sub-sample of these women who were between ages 10 and 30 in 1968. The adjusted annual household income for each of these women was constructed by using the regression results to predict income for each year in the absence of cyclical and cohort effects. To eliminate cohort effects, the individual fixed effects were adjusted to remove the effects of the women's age in 1968, while the cyclical effects were purged by omitting the estimated effect of the US unemployment rate each year. [9]

The resulting distributions of lifetime income are summarized by the columns of Table 1. The first two columns show that labor income is more unequally distributed than total income: the coefficient of variation, which measures dispersion, and the skewness are both higher for labor income. The higher Gini coefficient for labor income reflects both of these measures,

[9] This unemployment effect was estimated separately for different categories of race, birth year and education.

showing that labor income is more highly concentrated in the high-income households. The "Percentile" section of the table reports the income of nth percentile household, expressed as a fraction of the mean. Thus a household that is at the 10th percentile in the distribution of labor income earns only 30 % of the earnings of the median household. But in the distribution of total income, the 10th percentile household earns 38% of the total income of the median household.

The table also shows that the Gini coefficient is much smaller for total income than for labor income; this is due to the inclusion of transfer income in total income. Asset income is also included in total income, but this has a disequalising effect that is evidently smaller than the equalising effect of transfer income. The relative magnitudes of these effects may however be due to under-representation of the very wealthy in the PSID, and to the exclusion of pension wealth, as suggested by Gustman et al (1999).

Table 2: Intergenerational Correlations of Lifetime Income*

	Total Income		Labor Income		Adjusted Labor Income.	
Mom's Income	0.43 (0.000)	23.05 (0.428)	0.38 (0.000)	-11.35 (0.158)	0.38 (0.000)	-8.18 (0.250)
Mom's Income Squared		-1.66 (0.474)		1.06 (0.128)		0.80 (0.200)
Mom's Income Cubed		0.04 (0.514)		-0.03 (0.114)		-0.02 (0.176)
R-square	0.13	0.15	0.14	0.15	0.16	0.17
N	327	327	327	327	327	327

**Lifetime income measures are based on author's computations from PSID 1967-1991. Sample composed of daughters aged 10-20 in 1968 and their mothers.*

To estimate the intergenerational persistence of lifetime income, a subsample of the income dynamics sample was constructed by identifying the population of daughters as all women who were aged 10 to 20 years old in 1968 and who were reported as living with their parents at that time and who became household heads or spouses by 1987. Linking these women with their mothers resulted in a sample of 327 mother-daughter pairs. The intergenerational income correlation was defined as the coefficient of the log of mother's income in a regression equation where the dependent variable is the log of the daughter's income. Polynomial terms in the age of the mother and the daughter are also included, in order to control for any age effects that may remain in the estimate of lifetime income. This procedure is quite standard; apart from the measure of lifetime income, it is identical to that described by Solon (1992) and Mulligan (1993).

The estimated correlations between mother and daughter are shown in Table 2, with probabilities of the t-scores underneath the coefficient estimates. The correlations are well within the range reported in the literature for the correlation of average income across generations, but somewhat lower than those reported in Knowles (1999a) for lifetime income. The correlation of total income is 0.43 which is somewhat higher than that of labor income, at 0.38. Including polynomial terms in mother's income did not change very much the explanatory power of the regression, suggesting that the linear specification is correct. The fact that the correlation appears lower for the women than for men can also be interpreted as the effect of more uncertainty in our estimate of lifetime income, as discussed above.

Table 3: Fertility and Transfers*

Earnings Percentile	Number of Kids	Transfer Share of Income
1	3.11	0.81
5	2.83	0.52
20	2.67	0.19
40	2.53	0.09
60	2.45	0.06
80	2.53	0.03
90	2.33	0.02
95	2.27	0.01
99	2.14	0.00

**Author's calculations from PSID representative sample of women aged 10-30 in 1968 (N=884). Earnings are author's estimate of lifetime household labor income.*

To estimate the fertility-earnings relationship, it is necessary to account for the variation across cohorts in the mean number of children. Where the younger women are concerned there is another problem: at the time of their most recent appearance in the PSID, they may not have completed their fertility. The way these problems are dealt with here is to estimate an equation relating the number of children to adjusted earnings and age, allowing for non-linear effects by including polynomial terms for each variable. The sample for these estimates is the sample of mothers from the intergenerational estimation above. The results are shown in Table 3. The table shows that the average number of kids by earnings percentile declines rapidly with earnings over the poorest half of the earnings distribution. The poorest women had an average of 3.11 children, while women at the 90th percentile had 2.14 children on average.

An important component of income in the model is transfer income. In the data this is defined to include income from government welfare programs, food coupons and transfers from relatives. The data suggests that the poorest women in the sample get most of their lifetime income from such transfers. However it was not possible to follow the above procedure for estimating lifetime transfers, as the income dynamics equation had very low explanatory power. Instead the final column of Table 3 reports the average percentage of a woman's income between ages 30 and 53 that was accounted for by transfers. Only women who had at least 5 observations in this age range were included. The results show that the women in the lowest earnings percentile receive an average of 81% of their income from transfers over this 25 year period, and that the importance of transfers declines rapidly with earnings over the lower half of the earnings distribution.

3.2 Model Specification

In computing the model, standard functional forms are used for the preferences, the technology and the stochastic process; the critical decision is the choice of parameters for these functions. Although some of these parameters are restricted by the results of previous research, most are not directly observable. This latter class of parameters is chosen so as to generate a good fit of the model to the U.S. data, and then the model is evaluated by the plausibility of these parameter values.

While most papers with endogenous fertility allow family size to be chosen on $\mathbb{R}_+$, it is both computationally convenient and realistic to restrict family size to a small set of discrete choices. In addition, the maximum number of children $\overline{n}$ limits the choice of the time cost parameter θ. The fertility choice set in the benchmark model was the discrete set $N = \{0, 0.5, 1, 1.5, 2, 3, 4, 5\}$. This allows for fractional number of children because of the possibility of a pair of real-llife parents choosing odd numbers of children, an option that is more important when choosing a low level of fertility than for higher higher fertility.

Parameterizing the preferences involves choosing a utility function for consumption and an altruistic discount function. Parental utility for consumption is given by a standard CES utility function specification: $u(c) = \frac{1}{\sigma}c^{\sigma}$. The discount factor function is given by $b(n) = \beta n^{-\gamma}$, where the parameters satisfy $\beta \in (0,1)$, and $\gamma \in (0,1)$; this is similar to Becker and Barro (1988), and Tamura (1994). As σ declines to zero, increasing the curvature of the utility function, the more attractive is increased fertility as a response to increased wealth. However the curvature γ of the discount factor reduces the benefits of increased fertility.

The discount factor β is set to be consistent with previous studies of wealth distribution, where typical values range from an annual discount factor of 0.91, as in Huggett (1994), to 0.96, as in Aiyagari (1994). Compounded annually over one generation (26 years), these values imply a range of 0.09 to 0.35. The discount factor parameter β in the model was set to 0.3 so that for one-children families[10].

The technology of the model consists of the child-raising costs, the stochastic function relating parental investment to children's wages, and the aggregate production function. This paper uses the standard Cobb-Douglas speci-

[10] An empirical counterpart of the discount function can be inferred from the differences in discount factors for child quality across income and family-size groups, as estimated by Agee and Crocker (1996) in a recent study of parental investment in the health of their children. They find that parent's discount rates for child quality are 3.2% annually for rich parents; if these are families with one child per parent, then this implies an intergenerational discount factor of 0.43; if these parents have more children, then the discount factor is even larger.They found that parents below the mean income level had an annual discount rate (6.8%) for child quality more than twice that of the parents above the mean. This implies a huge difference in inter-generational discount factors: 0.16 for the poor, 0.43 for the rich, over a 26-year span. Assuming that the rich have two children and the poor four (average for the sample is 2.9), this means that $2^{-\gamma} = 0.16/0.43$, so $\gamma = 1.6$ is the implication of this approach(Note that the model's equivalent of n children per family is $n/2$ children per parent.).If borrowing constraints affect the decisions of the poor in the empirical study, it may be that their discount factors were underestimated, which would mean that the curvature parameter is overestimated. This is far more curvature than implied by the benchmark calibration, so the current model errs on the side of conserving the structure of standard models.

fication for the aggregate production function, $F(K, L) = K^{1-\alpha} L^{\alpha}$, with the labor-share parameter $\alpha = 0.67$.

The time costs of children are very difficult to estimate empirically; standard procedure in the empirical literature is to use an instrumental variable like mother's education as a proxy for the child-raising costs, as in Mulligan (1993). Attempts to measure the impact of children on mother's labor supply include Angrist and Evans (1996) and Rosenzweig and Wolpin (1980). The former of these resulted in an estimate of the time cost per child of roughly 17% of the mother's labor time while the child was young. This is likely to be an underestimate of the time cost of children, since it was estimate for the margin between 2 and 3 children, whereas the costs are likely to be much larger for the margin between 0 and 1. Even a generous padding of this estimate still results in a small time cost in terms of the model. Assuming a work career of 40 years, and that the child cost is 10% of a couple's annual labor endowment for 10 years results in a total time cost per child of 2.5% of the time endowment. This number may still be too low because the estimated time cost does not include the effect of children on parental leisure, which is likely to be just as large, given the lower rate of employment of married women.

Rosenzweig and Wolpin (1980) use the labor-participation decisions of women whose first children are twins to identify the labor-force participation effects of an additional child. Thus they measure the margin between 1 and 2 children, which is likely to be larger than that between 2-3 and smaller than that between 0-1 children. They find that having twins as a first birth lowers the probability of the mother's current labor-force participation by about 0.22 for mothers aged between 15 and 34. This effect confounds timing effects with family-size; women with twins actually spend more time in the labor force over the lifecycle

Aggregate measures of child costs can be inferred from Haveman and Wolfe (1995), who computed an upper bound for time costs of children as the reduction in mother's work time from full-time work, valued at the average wage for women of the same level of education as the mother. The time cost per child according to this procedure was $4,000 per year in 1992. Assuming an average female full-time salary of $20,000 in 1992, the per-child time cost is 1/10th of the total parental time endowment. This measure excludes the effects of children on the leisure of both parents, which is likely to be non-negligeable, though problems may arise in distinguishing child-care time from leisure. Assuming that the leisure costs per parent equal half of the women's loss of labor time would imply a per child cost of 20% of the parental time endowment. A reasonable range for the time cost parameter θ then would be somewhere between 5% and 10%.

The per-child goods costs of child-raising are given by the function $\phi(e', s) = \left[e'/(e^{\zeta})\right]^{\eta} + \lambda$, where $\eta \in (1, \infty), \zeta \in [0, 1)$ and $\lambda \geq 0$. This was chosen for simplicity; costs are convex, and independent of the parent's luck, though not of her education. The effect of additional parental investment on children's education is determined by two parameters: the curvature of the education cost function with respect to children's education is given by η, and the dependence of the production function for children's education on the education of the parent by the parameter ζ. Increasing η increases the costs of higher levels of education relative to the costs of lower levels. Therefore inequality in education will tend to fall as η increases. The greater the curvature of the education cost function, the greater the benefits from spreading out investment across children for the families whose education choice remains constrained. Poor people will therefore tend to have larger families as η increases. Unconstrained parents will respond to higher income by increasing bequests, so their fertility will be less affected by an increase in η. The parameter ζ increases the education of the children of more educated parents, and so also contributes to inequality, but more specifically, to the intergenerational correlation of earnings among the more educated parents[11].

The function relating adult human capital h to education and market luck is $h(e, z) = e(1 + z)$. Market luck shocks z_t are assumed to follow the log-normal Markov process:

$$\log z_t = \rho \log z_{t-1} + \sigma_z \varepsilon_t \quad \varepsilon_t \sim N(0, 1)$$

Parents with higher levels of the shock have a higher opportunity cost of time and therefore have fewer children. As σ_z increases, so does the skewness of the shocks distribution. A positive persistence parameter ρ lowers the expected fertility of the children of parents with good labor luck, which increases the benefit of educating the children, and also directly increases the marginal benefits of education. Therefore a small increase in persistence causes quite large increments in the optimal education policy of lucky parents, and consequently, significant reductions in their children's fertility. This

[11] In this specification, the derivative $\frac{\partial^2 \phi}{\partial e^2} > 0$ ensures that the value function is concave in education, so that optimal education choices are bounded above.

process is approximated by a 7-state Markov chain using method described in Tauchen (1986) (see Appendix for details).

The remaining parameters are set so that the model's steady-state matches features observed in the data in the empirical section of the paper. These parameters are those associated with the transfer $\delta(y)$, the education cost function ϕ, the consumption-elasticity of utility σ and the process for the shocks, z_t. The features in the data that are used to restrict these are:

- The earnings-income relationship: lifetime earnings increases with lifetime income,
- The fertility-earnings relation: fertility declines with lifetime earnings,
- The transfers-earnings relation: transfers as a share of lifetime income decline with lifetime earnings,
- The inter-generational correlation of earnings,
- The earnings distribution: Gini coefficient and log-normality.

Matching these features in the model economy does not guarantee a unique set of 'best' parameters. However replicating these conditions is a non-trivial excercise, and because the model is non-linear, there is no guarantee that any given specification can achieve this. For instance rich parents invest less in their children's human capital and provide them instead with bequests; both features tend to reduce the earnings of the rich and increase their fertility, so the first two items are often violated for otherwise plausible choices of parameters. It is easiest to resolve this by increasing the variance of the shocks process, but this is limited by the need to match the Gini coefficient for earnings. In the end, the goal is to match the features of earnings inequality and fertility that are most relevant to parental decisions in the model.

3.3 Solution Method

The solution method to find the steady-state of the model is an extension of methods described in Huggett (1994), Aiyagari (1994) and Rios-Rull (1995). The method consists of solving the economy given a conjectured capital-labor ratio, computing the implied steady-state capital-labor ratio, and updating the conjecture. The procedure is repeated until the guess for the capital-labor ratio is equal to the value implied by economy's steady-state distribution. Applying this procedure requires computation of the parental decision rules, given the capital-labor ratio, and the steady-state distribution, given the decision rules.

The parental decision rules are found by solving the parental problem, given the interest rate and wages implied by the capital-labor ratio. Given these prices, the parental problem is solved by value-function iteration over a 3-dimensional state-space grid: starting with a conjectured value function, the optimal decisions are found for each point (e, a, z) on the state-space grid by comparing the parent's welfare from each feasible family size n. Given n, the algorithm uses the first-order conditions for investment to calculate the optimal allocation of income over parental consumption choice c and investment per child (e', a'). Then these solutions are used to update the value function, and the method repeated until the value function converges.

The steady-state distribution of agents over states implied by the resulting decision rules is found, given an arbitrary initial distribution. The procedure here is to consider the operator T implied by the decision rules and the transition matrix for the shocks. This operator maps distributions of parents over the state-space back into the space of distributions; given sufficient mixing among the states, the mapping has a fixed point. The procedure in this paper computes this fixed point by repeatedly applying the mapping to the initial distribution while keeping track of all the states with positive mass. The distributions of education and savings usually converge within 25 iterations to a very high degree of precision; for the benchmark economy, the capital-labor ratio, the standard-deviation and skewness of labor and savings all converged under the sup norm from a variety of different initial distributions.

The uniqueness of the steady state arises because the specification of education costs and the limit on family size ensure that the poorest parents educate their children sufficiently so that, with good market luck, they will earn more than their parents. Therefore even the poorest families are not trapped into certain poverty for all future generations. On the other hand, extremely high wealth in the model may be an absorbing state: children from the very richest families do not spend much time on the labor market, and so are less susceptible to poverty arising from bad luck. This analysis is shown in Figure 3. In each picture, education or income of descendants of the richest household in the state space is given by the line that starts from top of the vertical axis, that of the poorest by the line starting from the bottom. The horizontal axis plots the number of generations elapsed. It is apparent from the diagram that the income and education of the wealthiest families will decline over time if the family is hit with a long sequence of the worst shocks to labor-market luck. More important however, is that this steady state lies below that of the poorest possible family hit by a long sequence of the best possible shock.

Since we have both that the rich dynasties eventually become poor and that the poor dynasties eventually become rich, both with positive probability, this global mixing ensures that the steady-state distribution is unique. Nevertheless, it is possible to parameterize the model so that the poor dynasties never become rich; in such a case multiple steady-states would result. However such a parameterizations were not consistent with the features of the data that the model was required to match, because they imply too high an inter-generational correlation of earnings among the poor.

4 Results: The Benchmark Model

The benchmark economy is the steady state equilibrium resulting from a parameterisation of the model that satisfies the conditions discussed above. The parameters of this economy are shown in the appendix in Table A.4. Average fertility in the steady state of this model economy is 1.17 children per parent, compared to an average of 1.12 in the PSID for cohort of women aged 10-30 in 1968. Figure 4 shows the average fertility by earnings percentile in the model's benchmark population. It is apparent that the earnings-fertility relationship in the model is also similar to the empirical relationship found in the PSID. In the model, as in the data, fertility decreases in earnings over the poorest 60 percent of the population[12], increases slightly over the fourth quintile, and declines further in the top decile.

The model succeeds in matching some key measures of income inequality and intergenerational persistence. Consider the measures of inequality of the empirical earnings distribution that were shown in Table 1; some analogous results for the model are given in Table 5. Comparison of these tables shows

[12]Since the dataset on which these observations are based under-represents the wealthiest Americans, the output of the model does not show the wealthiest 10% of the steady-state distribution.

that the Gini coefficient of the model's earnings distributions (0.35) is close to that observed in the PSID (0.34). The resulting earnings distribution is also shown in Figure 4, which shows a roughly log-normal shape for the model, as in the data. The inter-generational earnings correlation is 0.38, but for total household income, the correlation is much higher, around 0.78. This reflects the role of wealth in the model, which is discussed below. For now note that while this estimate is much higher than the estimates in Solon (1992), it is close to the father-son correlation of 0.71 estimated by Mulligan (1993) using instruments for parental income.

Table 4: Statistics for Model Economies*

	Bench-Mark	Constant Fertility
Mean Income	1.00	1.45
Median Income	0.85	1.40
Earnings Gini	0.35	0.34
Income Gini	0.39	0.35
Wealth Gini	0.91	0.72
Pct. Low Educ	0.17	0.02
Pct. $<$ EducBnch/2	0.17	0.02
IG Corr Earnings	0.38	0.34
IG Corr Wealth	0.30	0.31
Av. Fertility	1.17	1.15
Savings Rate	0.03	0.06
Capital-Labor Ratio	2.59	2.65

**Computed from steady-state equilibria of benchmark and constant-fertility economies. Gini coefficients for earnings and income exclude richest 5% of population.*

Given the success of the model in matching these features of fertility and earnings inequality, it is necessary to evaluate the parameter values that were required to obtain this benchmark model. The value for the utility curvature is the parameter the model is most sensitive to, because it governs the trade-off between quality and quantity. A recent study by Keane and Wolpin estimated this parameter at around 0.35, using the curvature of consumption profiles of middle-aged men. This is quite close to the benchmark value. [13]

[13] Keane and Wolpin argue that the standard estimates of this parameter are much lower,

It turns out that the model is also quite sensitive to the time-cost parameter, θ and the transfer, δ. The first of these was set to 0.08, which is well within the range deemed reasonable by the arguments of the preceding section. This value implies that a parent may have the maximum of 5 children, and still supply 60% of her total time endowment to the labor market, which is analogous to a two-parent family having 10 children, and one spouse working full-time, while the other works half-time over 20% of her adult life. Figure 4 also shows that the parameterisation of the transfer δ resulted in a distribution of transfer income as a share of total income that understates the cross-sectional pattern in the data, except for the 20th income percentile which it matches exactly.

According to Haveman and Wolfe (1995), total spending on children's human capital in the in the U.S. is 14.5% of total output. In the model, parents spend 12% of total output on children. In Table **??**, the statistic 'Pct of Kids with Low Educ.' shows that the fraction of kids with less than half of the efficient level of education is 16% in the benchmark model. [14] This statistic suggests that under-education due to binding parental budget constraints is a significant problem in the benchmark economy.

4.1 Wealth Accumulation in the Benchmark Model

These results suggest that the parental decisions model is indeed quantitatively consistent with lifetime income inequality in the U.S. What then does the model imply about wealth? Wealth holdings in the model are defined as total savings for bequest purposes. The model's steady-state was solved under the assumption that this bequest motive was the only source of wealth accumulation; the capital-labor ratio implied by this parameterization may be significantly higher if other motives, such as precautionary or life-cycle considerations led to additional wealth accumulation. More relevant to the current model however, the savings level in the model would be higher than

because previous research has identified it from the curvature of the consumption profiles of young men, implictly assuming away borrowing constraints that limit the ability of young men to transfer consumption forward in time.

[14] Given the parental state, this efficient level is defined as the maximum education level chosen by parents with the same level of market luck and education. Persistence in the market-luck process makes the efficient level higher for parents with high realisations. High levels of parental education reduce the cost of education per child.

the bequest level to the extent that human capital investments also enter savings (e.g. parents save for college), as argued by Kotlikoff and Summers (1981). In the model, because the entire lifespan is one period, spending on children's human capital does not enter savings at all, unlike in reality, where at least the portion that is saved in anticipation of children's college expenses is a significant part of total savings for middle-class families.

Table 6 compares the shares of total wealth held by households in the benchmark economy to the findings of some recent empirical studies of wealth inequality, Diaz-Gimenez, Quadrini, and Rios-Rull (1997), and Wolff (1994). These studies differ in their definition of wealth, but both attempt to include household durables, owner-occupied housing and other categories of non-financial wealth. The benchmark model overstates the data in that all of total wealth is concentrated in the richest quintile. The Gini coefficient for wealth in the benchmark model is 0.90, while according to the empirical studies cited in Table 6, the U.S. wealth distribution has a gini coefficient between 0.72 and 0.78. The model understates the wealth of the richest 1% of households, who hold 23% of wealth in the model, compared to nearly 30% in the data. The low wealth holdings of the poor in the data, which are difficult to explain given other savings motives, are explained here by the higher rate of return, for constrained parents, on children's human capital.

Table 5: Percent Share of Total Wealth*

Source	Wealth Quintiles		Wealthiest Percentiles			Gini
	4th	5th	10 to 5%	5 to 1%	top 1%	Coefficient
Diaz-Gimenez et al (1997)	13.43	79.49	12.62	23.95	29.55	0.78
Wolff (1987)	6.00	75.00	N/A	21.00	28.00	0.72
Benchmark Model	0.00	100.00	27.61	36.34	22.88	0.90
Constant-Fertility Model	19.82	79.70	21.60	19.78	8.01	0.76

**Wealth in model economies is computed as total savings for bequests in steady-state equilibrium.*

If savings are strictly interpreted as bequests, then the model does not yield a very high rate of savings (about 3%), so these results do not support the contention of Kotlikoff and Summers (1981) that inter-generational transfers play a major role in U.S. savings as a whole. However total spending on children in the model is 23% of income, so to the extent that U.S. households save in order to invest in their children, as Kotlikoff and Summers (1981) argue, then the inter-generational motive may indeed play a large role in explaining savings.

To summarize, the basic result of this exercise is that the parental decisions model is a plausible model of U.S. inequality: the model generates an equilibrium steady-state earnings distributions that turns out to be consistent with the most-used measures of U.S. earnings inequality. This distribution was obtained under the restrictions that the fertility-earnings relation generated by the model resemble that of the data, and that parameter values, such as the time cost of children, be within the ranges suggested by previous research. Although this model fails to explain the extent of wealth concentration in the richest 1% of U.S. households, the model does much better in this regard than the standard general-equilibrium models of precautionary and lifecycle savings.

4.2 Does Fertility Matter?

How different would the U.S. wealth distribution be if fertility were the same for all parents? If this margin does not make a significant difference in average income, inequality, or absolute poverty rates, then it may well be reasonable for macroeconomic models to ignore it altogether, and concentrate instead on the other trade-offs parents face. To evaluate the role of fertility, the benchmark model was modified by fixing fertility per parent at the benchmark average for all parents and re-computing the model's equilibrium steady state, subject to a stochastic process that is recalibrated to match the earnings inequality of the benchmark economy.

The aggregate measures from the resulting steady-state equilibrium are shown in the 'Constant Fertility' column of Table 5. With the capital-labor ratio held constant, eliminating the fertility margin results in a slight increase in per capita income, and a large decline in wealth inequality. Median income increases by much more than mean income, suggesting that income inequality is much lower with constant fertility; the Gini coefficient is now 0.356, compared to 0.39 for the benchmark. The intergenerational correlation of wealth has not changed much, but that of earnings has fallen from 0.38 to 0.34.

The proportion of parents who choose less than half of the efficient level of human capital investment for their children is 2% in this economy, compared to 16.7 % in the benchmark model[15]. In other words, output per capita

[15] The efficient level is computed as the maximum level chosen by a parent with the same education and the same realisation of market luck.

would increase significantly if it were possible to increase the human capital of the children of poor parents without distorting parental decisions. This does not imply, however, that even this poorest 17% of parents would significantly increase investment in their children's human capital if they were to receive transfers, since they might also respond by having more children or consuming more. The point is that the quality-quantity trade-off plays an important role in explaining educational inequality.

From Table 5, it is clear that the average savings rate is much higher when fertility is fixed. For the question at hand, the main result of suppressing fertility decisions is that the wealth distribution is becomes more compressed: the Gini coefficient falls from 0.90 to 0.76. Table 6 shows that wealth holding in the 4th quintile is now significant, and that the concentration of wealth in the richest 1% has fallen from 23% in the benchmark model to only 8% of total wealth. Thus endogenous fertility makes a strong contribution to explaining this aspect of the wealth puzzle.

To summarise, suppressing fertility variation in the benchmark model results in a significant reduction in both inequality in wealth and the persistence of labor income inequality across generations. This suggests that endogenous fertility might play a key role in understanding the U.S. wealth distribution.

5 Conclusion

This paper set out to assess a simple explanation of two well-known anomalies in the U.S wealth distribution: the concentration of wealth in the richest 1% of households, and the low wealth holdings of most of the rest of the population. The explanation relied on the quality-quantity trade-off in parental investment decisions, which was embedded in a general-equilibrium model of the earnings distribution. The logic of the exercise was that if this the quality-quantity trade-off were indeed a plausible explanation of U.S. wealth inequality, then a version of the model calibrated to match U.S. earnings inequality and the fertility-earnings relationship would also generate a wealth distribution with the anomalous features of the U.S. distribution. If fertility were indeed the key component of this explanation, then suppressing fertility variation in the model should hinder the model from generating these anomalies.

To carry out this procedure, this paper constructed and solved a model of the income distribution based on a simple theory of parental fertility and investment decisions. Since the model represented the adult life-span by one period, it was also necessary to estimate distributions of life-time earnings and income from the PSID and compute relationships between age-adjusted fertility and earnings, and between earnings and the fraction of income received as transfers. The model was calibrated by choosing standard functional forms and parameter values that fit the range of plausible values established by previous research. In the cases where previous research did not restrict the choice of parameter values, these were restricted by requiring the model to fit the observed relationship between fertility and earnings.

This benchmark version of the model did in fact generate the two anomalies in question. In fact, wealth holding of the poorest 80% of households was lower than in the data. Although the holding of wealth of the richest 1% of households was less than in the data, it was much higher than was the case for the standard models reviewed in Diaz-Gimenez, Quadrini, and Rios-Rull (1997). Furthermore, when fertility variation was suppressed, wealth inequality in the model economy fell significantly; in particulark, the concentration of wealth in the richest 1% of households fell from 22% to 8%. From this one can infer that the success of the model with regards to wealth inequality is due to the importance of the human-capital and fertility margins for the investment decisions of most households; these margins are generally omitted in both theoretical and empirical studies of wealth accumulation.

Overall, these results would seem to support the idea that fertility plays a key role both in explaining wealth inequality and transmission of earnings inequality across generations. However an important caveat is that the analysis abstracted from other savings motives. Indeed given that the results showed a rather low savings rate, relative to the actual savings rates observed in the U.S. until recently, it would be interesting to pursue this research further in a richer environment that allowed a role for precautionary and lifecycle savings motives, as well as fertility and human capital investment.

References

Agee, M. D., and T. D. Crocker (1996): "Parent's Discount Rates for Child Quality," *Southern Economic Journal*, 63, 36–50.

Aiyagari, S. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 659–684.

Alvarez, F. (1994): "Two Essays on the Economic Theory of Fertility," Ph.D. thesis.

Angrist, J. D., and W. N. Evans (1996): "Children and their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size," NBER Working Paper 5778.

Becker, G. S., and R. J. Barro (1988): "A Reformulation of the Economic Theory of Fertility," *Quarterly Journal of Economics*, CII(1), 1–25.

Becker, G. S., and N. Tomes (1976): "Child Endowments and the Quantity and Quality of Children," *Journal of Political Economy*, 84(4, pt. 2), S143–S162.

——— (1986): "Human Capital and the Rise and Fall of Families," *Journal of Labor Economics*, 4(3, pt. 2).

Castaneda, A., J. Diaz-Gimenez, and J.-V. Rios-Rull (1998): "Earnings and Wealth Inequality and Income Taxation," Mimeo, University of Pennsylvania.

Diaz-Gimenez, J., V. Quadrini, and J.-V. Rios-Rull (1997): "Dimensions of Inequality: facts on the US Distributions of Earnings, Income and Wealth," *FRB Minneapolis Quarterly Review*, 21(2), 3–21.

Haveman, R., and B. Wolfe (1995): "The Determinants of Children's Attainments: A Review of Methods and Findings," *Journal of Economic Literature*, 33, 1829–1878.

Huggett, M. (1994): "Wealth Distribution in Life Cycle Economies," Mimeo, University of Illinois.

Knowles, J. (1999): "Measuring Lifetime Inequality," PhD Dissertation, University of Rochester.

Kotlikoff, L., and L. Summers (1981): "The Role of Intergenerational Transfers in Aggregate Capital Formation," *Journal of Political Economy*, 89, 706–32.

Krusell, P., and T. Smith (1996): "Income and Wealth Heterogeneity in the Macroeconomy," Manuscript,University of Rochester.

Laitner, J. (1992): "Random Earnings Differences, Lifetime Liquidity Constraints, and Altruistic InterGenerational Transfers," 58, 135–170.

Loury, G. C. (1981): "Intergenerational Transfers and the Distribution of Earnings," *Econometrica*, 49(4), 843–867.

Mulligan, C. B. (1993): "On Intergenerational Altruism," Ph.D. thesis, University of Chicago.

Rios-Rull, J.-V. (1995): "Models with Heterogeneous Agents," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley. Princeton University Press, Princeton, New Jersey.

Rosenzweig, Mark, R., and K. I. Wolpin (1980): "Testing the Quantity-Quality Fertility Model: The Use of Twins as a Natural Experiment," *Econometrica*, 48, 227–240.

Solon, G. (1992): "Intergenerational Income Mobility in the United States," *American Economic Review*, 82(3), 393–406.

Tamura, R. (1994): "Fertility, Human Capital and the Wealth of Families," *Economic Theory*.

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181.

Wolff, E. N. (1994): "Trends in Household Wealth in the United States: 1962:1983 and 1983:1989," Economic research report, C.V. Starr Center for Applied Economics.

Zimmerman, D. J. (1992): "Regression Toward Mediocrity in Economic Stature," *American Economic Review*, 82(3), 409–429.

A Appendix

A.1 Estimation Results for PSID

The following table shows the characteristics of the income dynamics sample of women described in the text.

Table A1: Earnings Sample Descriptive Statistics*

Variable	Mean	Std Dev	Minimum	Maximum
BLACK	0.09	0.29	0.00	1.00
PRIMARY	0.98	0.14	0.00	1.00
COLLEGE	0.31	0.46	0.00	1.00
HISKOOL	0.79	0.41	0.00	1.00
NUMKIDS	2.50	1.90	0.00	14.00
AGE68	32.99	16.37	10.00	79.00
TOTINC3	$10,081	$6,601	$0	$79,069
LABINC3	$8,372	$6,089	$0	$77,333

**PSID cross-section; sample restricted to women aged 10-79 in 1968 (N=2183). Income is reported in 1967 dollars. Race is based on that of head of family in 1968.*

The following table shows the results from the estimation of an OLS equation with individual fixed effects. The dependent variables were annual levels of total household income and household labor income, transformed into logs. All of the variables below are interacted with years of potential labor experience. Coefficient estimates are generally significant at the .01 level, except for College and HiSkool in the total income regression.

These results were used to compute adjusted annual income (purged of cohort, cyclical and lifecycle effects. The present value of this adjusted income served as the basis for the lifetime income results presented in Table 1.

Table A2: OLS Parameter Estimates of Income Model*

Variable	Total Income			Labor Income		
	Parameter Estimate	Standard Error	t Value	Parameter Estimate	Standard Error	t Value
UNEMP	-8.06E-03	(2.25E-03)	-3.58	-2.09E-02	(2.99E-03)	-6.97
COLLEGE	2.71E-03	(1.47E-03)	1.85	-1.64E-02	(1.95E-03)	-8.37
HISKOOL	-7.06E-04	(1.33E-03)	-0.53	-6.30E-03	(2.05E-03)	-3.06
PRIMARY	-1.69E-02	(2.88E-03)	-5.85	-1.45E-02	(4.88E-03)	-2.96
BLACK	3.09E-02	(6.53E-03)	4.73	9.39E-02	(1.05E-02)	8.94
WHITE	4.95E-02	(3.41E-03)	14.51	1.25E-01	(5.36E-03)	23.28
BLEXP2	-4.73E-04	(7.86E-05)	-6.02	-1.85E-03	(1.54E-04)	-12.05
WHEXP2	-6.80E-04	(2.38E-05)	-28.56	-2.29E-03	(4.15E-05)	-55.11

**Computed from 37,771 observations on women aged 30-79 in PSID 1967-1991. Variables interacted with years since school completion. Explanatory variables include indivdual fixed effects.*

A.2 The Stochastic Process Approximation

The process for luck z_t is:

$$\log z_t = \rho \log z_{t-1} + \sigma_\varepsilon \varepsilon_t \quad \varepsilon_t \sim N(0,1) \tag{15}$$

For any $z_t \in \mathbb{R}$ the conditional CDF is given by

$$F[z_t|z_i] = \Pr\left[\varepsilon_t < \frac{z_t - \rho z_i}{\sigma}\right] = \Phi\left(\frac{z_t - \rho z_i}{\sigma}\right) \tag{16}$$

As in Tauchen (1986), I approximate this Markov process by a Markov chain with n states z_i. Let $n_1 =$ the integer portion of $(n/2)+1$ and denote the unconditional standard deviation of z_t by $\sigma_y = \frac{\sigma_\varepsilon}{\sqrt{1-\rho^2}}$. I then choose the states of the Markov chain so that:

$$\log z_i = \begin{cases} -n_1 \sigma_y & i = 1 \\ l_{i-1} + \sigma_y & i > 1 \end{cases} \tag{17}$$

Consider $j \in \{2..N-1\}$ and let $w = \log z_j - \log z_{j-1}$.The transition probability from state i to state j is defined as:

$$\begin{aligned} p_{ij} &= \Pr\left(\log z_j - w/2 < \rho \log z_i + \sigma \varepsilon_t < \log z_j + w/2\right) & (18) \\ &= F\left(\frac{\log z_j + w/2 - \rho \log z_i}{\sigma}\right) - F\left(\frac{\log z_j - w/2 - \rho \log z_i}{\sigma}\right) & (19) \end{aligned}$$

A.3 Details of Solution Method

Solving for the steady-state income distribution implied by a set of parameters and a capital-labor ratio involves two basic steps: solving the decision problem given prices, and solving for the steady-state distribution implied by these decisions, given an arbitrary initial distribution.

The parental decisions were solved by the standard technique of value-function iteration on a state-space grid. This grid consists of 3 dimensions; two of these are defined by 25 evenly spaced levels of human capital or inheritance between zero and a maximum chosen to be higher than the maximum that would be selected under the optimal decision rules. The third dimension consists of 7 values of the labor-luck shock, selected according to the method of approximating a continuous log-normal Markov process described above.

There are some non-standard features of the computational problem: (1) one of the choice variables, fertility, enters the budget constraint multiplicatively with two of the other choice variables, children's human capital and bequests, and (2) the discount factor is endogenous. These problems are dealt with here by restricting fertility choice to a small finite set. The first difficulty is resolved by solving the parental problem taking fertility as given and simply choosing the fertility level that results in the highest expected value, while the second is eliminated by restricting the fertility set so that the maximum attainable discount factor is less than one. Given an arbitrary guess of the value function, the strategy therefore is to solve the parental problem at a given state-space point for each feasible level of fertility, update the value function given the utility levels implied by the solution at each state-space point, and repeat until the value function converges.

Given fertility, the parental problem is to choose consumption c, human capital investment e' and bequest a' subject to the budget constraint (5). Because the interaction between e' and a' may result in multiple local optima, the algorithm uses three different solution techniques to identify candidate solutions to this investment problem, and chooses the candidate that yields the highest value to the parent. In all of these techniques, the idea is to identify roots of the first-order conditions for investment and consumption. The first-order conditions are evaluated by linear interpolation of the decision rules from the previous iteration.The first technique assumes that bequests are zero, and checks for roots of the first-order condition for e'. If the the first-order condition for a' is non-positive at these roots, then each is root is a candidate solution. The second technique uses a Newton-type method to search for roots of the system of first-order conditions, on the assumption that both e' and a' are positive. Finally, the algorithm looks for roots of the consumption first-order condition by solving for roots of the first-order conditions for e' and a', taking c as given.

The steady-state distribution is found by iterating on the following procedure, starting from an arbitrary distribution of parents over endogenous states. First a score is computed for each parental state, based on the state's mass in the distribution, and the savings and human capital; the idea is to ensure that even parental states with very low mass are included if they account for a significant fraction of labor or capital. For each parental state whose score is above some threshold, the probability of each realization of the stochastic state z, and hence the mass of each state (e, a, z) is calculated. The optimal decisions are calculated, taking as given the value function found by the preceding procedure. These decisions determine the distribution of children over endogenous states, and the amount of labor supplied by the parent's generation. The capital-labor ratio is calculated using the parent's inheritance as the measure of capital. The procedure is repeated until the average levels of capital and labor, and their dispersions, converge.

A.4 Benchmark Parameters

Table A3: Benchmark Functional Forms and Parameters

Utility function	$u(c) = \frac{c^\sigma}{\sigma}$	$\sigma = 0.43$
Discount factor	$b(n) = \beta n^{-\gamma}$	$\beta = 0.275$ $\gamma = 0.55$
Per-child goods cost	$\phi(e'; e, z) = \left[e'/(e^\zeta)\right]^\eta + \lambda_1 + \lambda_2 y$	$\zeta = 0.10$ $\eta = 1.8$ $\lambda_1 = 75$ $\lambda_2 = .015$
Time cost per child	$\theta(e, z) = \theta$	$\theta = 0.10$
Luck process	$\log z_t = z_0 + \rho \log z_{t-1} + \sigma_z \varepsilon_t$	$z_0 =$ $\rho = 0.10$ $\sigma_z = 0.8$ $\varepsilon_t \sim N(0, 1)$
Production function	$F(K, L) = K^{1-\alpha} L^\alpha$	$\alpha = 0.67$
Income Tax Rate	τ	$\tau = .0175$
Consumption endowment	$\delta_0(y) = \max(0, \delta_1 + \delta_2 y + \delta_3 y^2)$	$\delta_1 = 3000$ $\delta_2 = -1.36$ $\delta_3 = 1.875 \times 1e^{-4}$

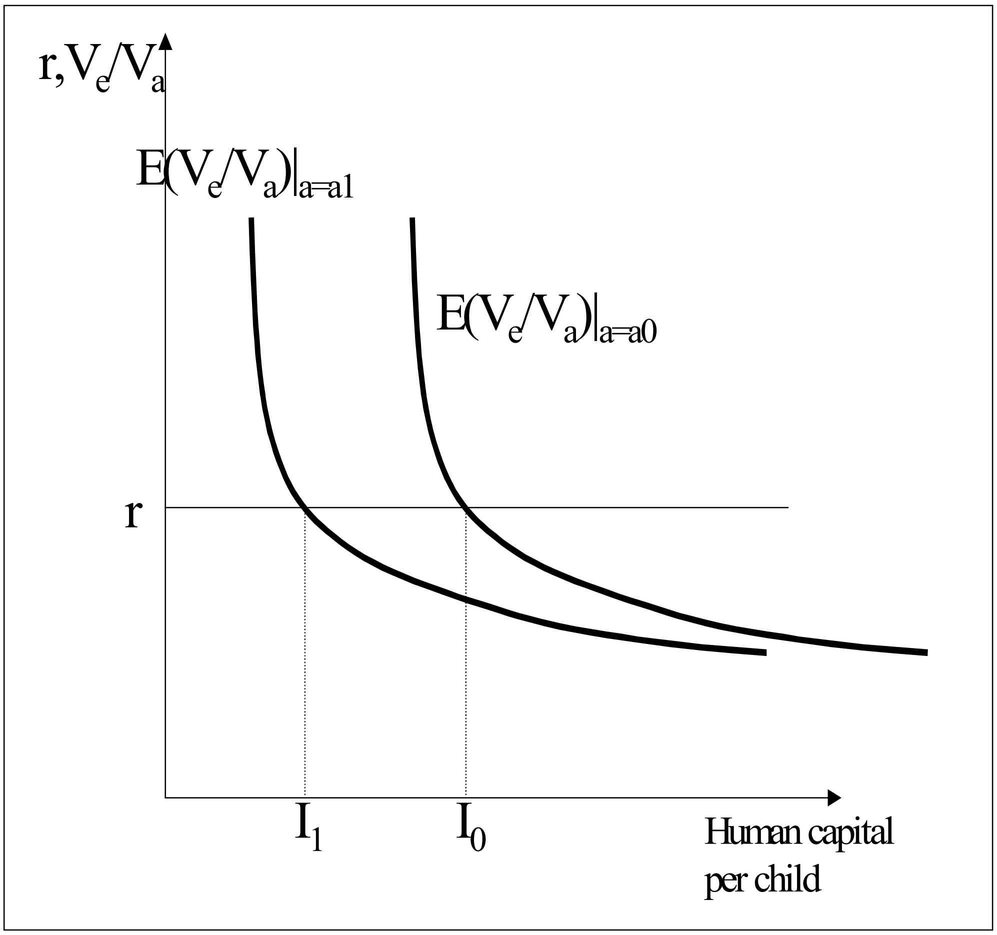

Figure 1: The trade-off between human capital and bequests. Given a bequest level of a_0,parents who invest less than I_0 per child have a higher return on human than physical capital. If $a_0 = 0$, then only parents who invest more than I_0 per child will make bequests. If $a_1 > a_0$, then return to human capital is lower, so optimal human capital falls.

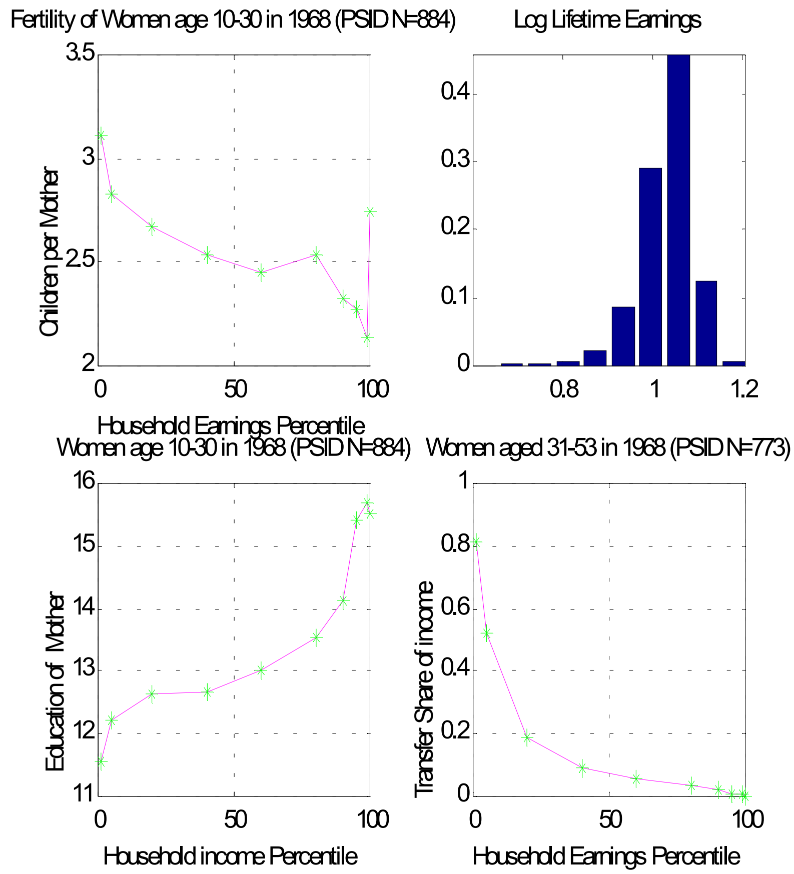

Figure 2: Lifetime income inequality in the U.S. Estimated by author from cohort of women aged 10-30 in 1968 PSID, except for transfer-income ratio, which is estimated from sample of women aged 31-53 in 1968.

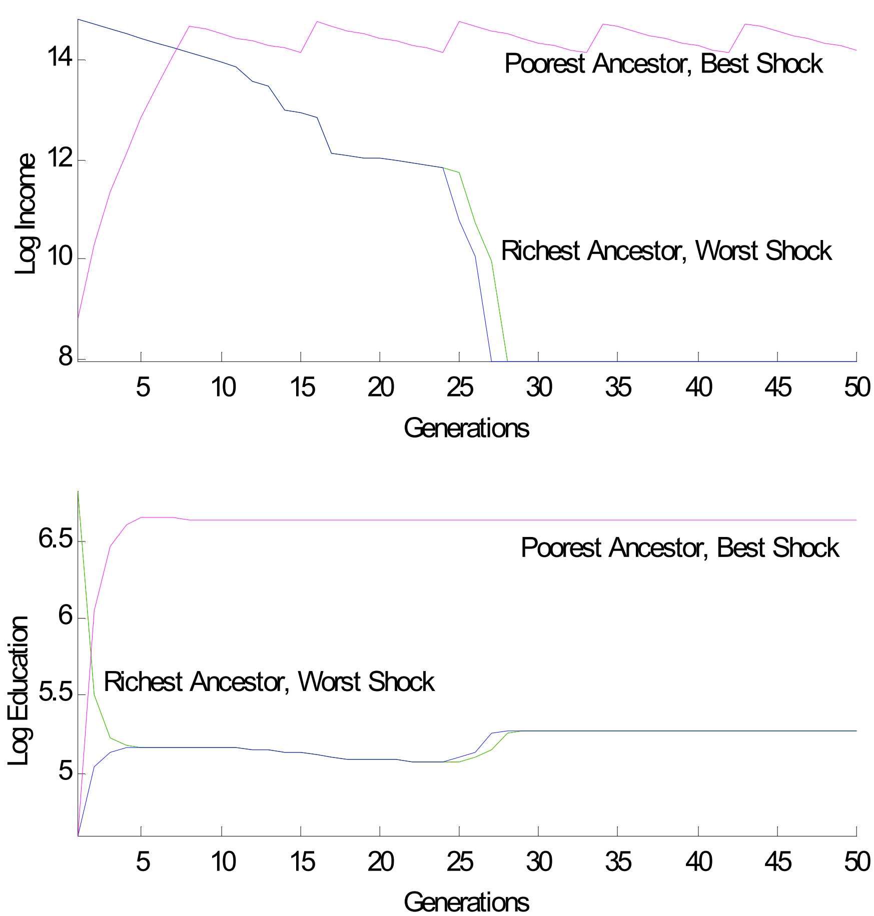

Figure 3: Education or income of descendants of the richest household is given by line that starts from top of vertical axis, that of the poorest by the line starting from the bottom.

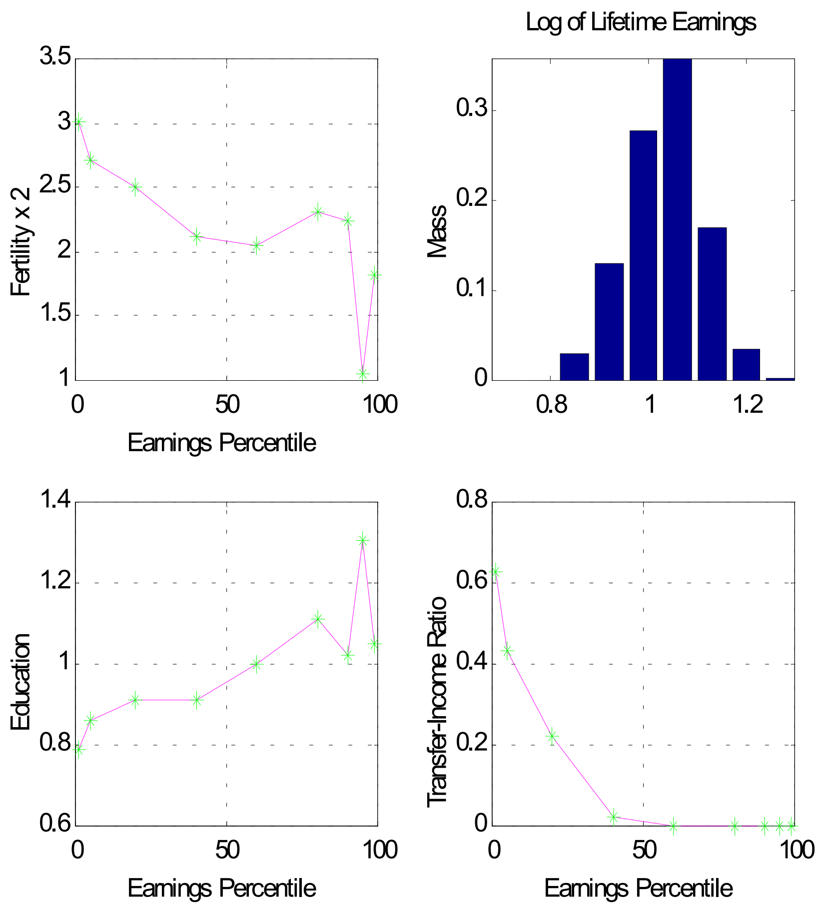

Figure 4: Inequality in the steady-state of the benchmark economy. Fertility is doubled to match 2-sex world. Education is reported as fraction of mean education.