Social Policy, Equilibrium Poverty and Investment in Children*

John Knowles†
Department of Economics,
University of Pennsylvania.

July 16, 1999

Abstract

Although it is well-known that redistributional social policies can have perverse long-run effects on the extent and incidence of poverty, it is not clear from the empirical literature how significant these effects are. This paper uses a dynamic general-equilibrium model calibrated to US household data to measure the changes in steady-state distributions of human and physical capital resulting from policy-induced distortions of parental decisions like fertility and investment in children. The results include measures of the net benefits by household income level of redistribution in the short run, and of the distribution of welfare across households in the long run. Holding constant the behaviour of children, redistribution makes the poor better off. However when parents anticipate the response of their children to redistribution, they invest less in their children's future, and the long-run outcome is that all households would be much better off without redistribution.

*Based on Chapter 3 of my PhD dissertation, prepared for submission to the Economics Dept. of the University of Rochester. I thank 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

Given the enormous cost of mistakes in social policy, it is natural that economists turn to empirical studies of existing policies to assess the probable effects of policies as yet untried. However Lucas (1976) provides a compelling reason for basing policy evaluation on equilibrium theory: prediction of the effects of new policies from past economic performance requires understanding explicitly how the previous behavior of the economy depended on the policies themselves. In macroeconomics, the result of the Lucas critique has been a new orthodoxy in policy evaluation that emphasizes structural modeling and dynamic equilibrium. In the study of social policy, however, this approach has been largely neglected until very recently, perhaps because of the notorious intractability of models of income inequality. This paper argues that such neglect may result in quantitatively important misapprehensions of the probable effects of redistribution, particularly regarding the long-run effectiveness of policy in reducing poverty or improving the lot of the poor.

For social policy, the problem is that the income distribution depends on the decisions of households, which in turn depend on the government policies; therefore the dominant effects of social policy in the long run might not be the direct effects measured by econometric studies, but the effects that only appear as the income distribution responds over time to the new policy regime. An example of this reasoning is implicit in the recent debate in the U.S. over welfare reform, some of which focused on the potential distortion in the fertility and child-rearing decisions occasioned by welfare programs. Although empirical research seemed until recently to suggest that these distortions were insignificant, according to Moffitt (1997) most recent studies find a positive effect of welfare transfers on fertility, and a negative effect on marriage. Moffitt acknowledges that there is still considerable disagreement among empirical researchers over the size of these effects, and that "a significant minority" of these studies still find no such effects, but concludes in favor of the existence of significant distortions. Under this hypothesis, a model of the interaction between the income distribution and parental decisions can be useful in evaluating the long-run effect of social policy.

There are at least three other reasons for taking a general-equilibrium approach to the effects of social policy on poverty. First is the interest of science: empirical studies of inequality across countries find that poverty rates are much higher in the U.S. . than in other industrialized countries; by comparing income and earnings inequality, Gottschalk and Smeeding (1997)

infer that differences in social policy constitute an empirically plausible explanation of these differences. An equilibrium model of inequality can tell us whether such policies do in fact generate the observed differences in inequality. A second reason is that getting social policy right may be one of the most effective ways of increasing per capita income available to the governments of industrialized countries, by raising the productivity of the large minority of people who currently earn far less than their country's median income. Finally, there is the interest of economic and social justice: equality of opportunity is a basic tenet of economic freedom, and hence a justification of social policy. An implicit assumption underlying all of these points is that social policy is effective in reducing poverty; this paper asks how true this is in the long run.

This paper contrasts the effectiveness of redistribution in general equilibrium with its apparent effectiveness when taking as given both the fertility decision rules and the distribution of labor income over households. The model used to effect these comparisons is a dynamic model of equilibrium income inequality and its transmission across generations through parental decisions. In this sense, and in the method to be used, this paper is related to recent general-equilibrium studies of welfare policy by Aiyagari, Greenwood, and Guner (1997) and of social security by Huggett and Ventura (1997). A fundamental difference between this paper and most of the previous literature however is that the current analysis incorporates accumulation of both human and physical capital[1].

The theory of the income distribution used in the current analysis is based on the benchmark model of Knowles (1999), which is calibrated to match the U.S. . income and fertility distributions. There are two basic features of that model that make it especially appropriate for this analysis: the theory is set in a dynamic general-equilibrium framework, and the earnings distribution is endogenous. In fact, the labor-income distribution in the model depends on precisely the sort of parental decisions that are thought to be susceptible to distortion by social policies, namely fertility and investment per child, but also allows, via a stochastic component of human capital, for earnings effects, such as luck in the labor market, that are outside the investment model.

If earnings of the children depend on the human capital investment de-

[1]Heckman, Lochner, and Taber (1998) also have this feature in their analysis of tax policy and human capital formation; however they focus on life-cycle decisions rather than inter-generational transmission of inequality.

cisions of the parents, and if these decisions are made subject to binding borrowing constraints, then redistribution has the potential to equalize the earnings distribution. This paper explores the long-run interaction between the earnings and wealth distribution on the one hand, and the returns to investing in human and physical capital on the other. The goal of the analysis is to determine how important quantitatively these distributional effects of social policy are likely to be in an economy similar to that of the U.S. The model used here will be augmented by the addition of a government that implements social policies which are funded by a proportional income tax. The method used is to simulate the U.S. income distribution using a dynamic general-equilibrium model calibrated to match emprical features of U.S. inequality.

By comparing the steady state income distributions that result from a prototypical redistributional policy, the analysis shows that the general-equilibrium effects of policy changes can be very potent, and large enough to actually reverse the equalizing direct effects of the policies, even when the short-run distortions appear minimal. There are two basic reasons why these general-equilibrium effects are important. First, small distortions in fertility and investment decisions are magnified by the way the distribution evolves over time. And second, the change in distributions imply changes in factor supply and hence relative prices of capital and labor, which in turn affect the investment decisions of all households, not just those directly affected by the policy.

The finding that redistributive taxation can be disequalizing in models of altruistic bequests dates back to Becker and Tomes (1979), although the theoretical insight itself is much older. The reason is that inheritances in these models are compensatory: parents save in order to insure children against bad realizations in the labor market. Therefore to the extent that the earnings of the children are independent of those of the parents, taxation increases inequality by disrupting this equalizing mechanism. Davies (1991) shows that when earnings are exogenous, taxing wealth redistribution can reduce long-run inequality, depending on the relative importance of earnings and non-labor income. Despite this long-run disequalizing effect of redistribution, however, Davies and Kuhn (1986) find that redistribution can be quite effective in the short run: in their model, inequality is likely to fall for several generations before approaching its new steady-state level from below.

The current paper differs from the above literature in several ways: (1) the concern of the paper is more with poverty than inequality, (2) the theory

underlying the analysis allows for decisions along margins that were omitted in previous analysis, such as fertility and investment in children's human capital, (3) the analysis combines a short-run partial equilibrium view with longer run general-equilibrium analysis, and (4) the analysis produces quantitative results that allow comparison of the size of the effects in question both among themselves and in terms of houshold income.

One of the difficulties in applying this intergenerational approach to evaluating social policies is that standard welfare comparisons are conceptually invalid: the average welfare of a household may be important, but how does this guide the comparison of policies that will result, via distortions of the fertiliy decision, in different numbers of households? Is a society less desirable because it has a higher number of households who are somewhat less well-off on average than a less-populous society? Even when the number of households is constant, can the greater happiness of an ancestor justify lower utility today for some households?

These are long-standing problems in welfare economics, and the response of this paper is to focus on a very specific type of comparison, by asking whether the poor households in the benchmark economy are better off than the poor in the steady-state economy that results from redistribution. The paper develops an explicit quantitative measure of the degree to which given households are better off in one economy than those in another, and offers these explicit comparisons in place of the inequality measures that are standard in the macroeconomics literature.

The main result of this paper is that the equilibrium effects discussed above are in fact strong enough, according to the model, to completely undermine in the long run any beneficial effect that the resdistributive scheme under consideration may have in the short run. Even when the fertility response to the social policy is suppressed, the effect on parental investment is sufficiently negative that an increasing fraction ends up with less human capital and hence higher fertility over time. Surprisingly, the principal effect of the redistributive scheme on parental decisions is not due to the direct incentive effects on the parents, but rather to the parent's anticipation of the children's response to the scheme. Since parents anticipate lower labor supply by the children, even rich parents invest less in the children's human capital. When fertility is flexible, poor households in addition choose to have more children, since the cost per child has fallen for two reasons: the marginal benefit of labor supply is lower, which lowers the time cost of children, and investment per child is lower, which lowers the cost per child.

The model does suggest that redistribution is very effective in alleviating poverty in the short run, even though the summary measures of inequality are relatively insensitive to the policy. The "short run" corresponds to the outcome of the policy when the distribution of households, and hence the wage and the interest rate are fixed as they were in the steady state without policy. Thus is not a true short run of the dynamic model, in the sense that it imposes myopia about price adjustments on the generation that is contemporaneous with the introduction of the policy. However this is close conceptually to the type of effect estimated in empirical analysis that attempts to infer the effects of social programs without fully identifiying the structural parameters of the behavioural equations, as discussed by **?**. In the experiments, the poorest half of the population experience significant increases in welfare over the benchmark case; the improvement, given a 5% tax rate on income, is equivalent for the poorest households to a 45% increase in income. Distortions of the fertility and the education decisions are shown to be quite small when the analysis abstracts from parent's anticipation of the effect of the policy on their children. Thus the model used here is consistent with empirical studies that find that changes in parental income have little effect on these decisions.

The following section presents the formal model that is used to derive these results. Section 3 presents the method for comparing economies. In Section 4, the analysis of the direct effects of redistribution is presented, and the steady-state analysis is presented in Section 5, which is followed by the conclusion.

2 The Model

The model is similar to that developed in Knowles (1999), except for the addition of a government sector that imposes a proportional tax on household income in order to fund lump-sum transfers to households. The model economy is composed of a continuum of agents who live for two periods, in an infinite series of overlapping generations. Each agent makes all of her decisions on entering the second period of life, i.e. on becoming an adult. Adult agents are distinguished by their education e , their inheritance a, and their luck in the labor market, z. They decide how many children n to have, how much to invest in education e' and how much to save for bequests a' for each child. Each decision is constrained to be non-negative; in particular, $a' \geq 0$

implies that parents cannot borrow against children's future income. The cost of children comprises both a reduction θ per child in labor-market time and a goods cost ϕ per child that is increasing in the education of the child, and decreasing in that of the parent, reflecting the idea that it is easier for more educated parents to educate their children to a high level of education. The market luck of the child is independent of that of the parent, and is not revealed until the child becomes an adult. Let the welfare of parents be given by U_0 and that of their children by U_1. Preferences over children and investment are given by the following recursive formulation, as in Becker and Barro (1988):

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

The transfer is modeled as an addition δ to the endowed income δ_0. The size of the transfers is determined by the revenue from the proportional tax at rate τ on all non-transfer income. The government budget constraint is:

$$\delta \square \tau \int [\delta_0 + y(e,a,z)]\, d\mu(e,a,z) \tag{2}$$

2.1 The Steady-State Equilibrium

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

$$y(e,a,z) = wh(z,e)\,(1 - g^n(e,a,z)\,\theta) + ra \tag{3}$$

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

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

subject to the household budget constraint:

$$\delta + \delta_0 + y(e,a,z)\,(1-\tau) - c - n[\phi(e') + a'] \geqslant 0 \tag{5}$$

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

The steady state equilibrium consists of a stationary probability measure $\mu(e,a,z)$, prices $\{w,r\}$, a capital-labor ratio $\left(\frac{K}{L}\right)$, and parental decision rules $g^x(e,a,z)$ for each choice variable $x \in \{n,e',a',c\}$, such that the parental decisions solve (4), the prices are implied by the capital-labor ratio, capital and labor equal the sum of the parent's inheritances and effective labor supply, respectively, and finally, the probability measure is consistent with the decision rules. These conditions are discussed more formally in Knowles (1999).

3 Issues in Evaluating Social Policy

The objective of this paper is to assess the importance for social policy of equilibrium interactions between parental decisions and the distributions of human capital and financial wealth. Specifically, the question is whether these interactions are likely to be important enough to undermine the ability of society to reduce poverty through redistribution. The method used to answer this question uses the benchmark version of the model presented in the previous paper to evaluate the effects on U.S. . inequality of a simple redistributive scheme.This approach raises a number of questions that are addressed in this section: what is the relevance of analysing an arbitrary policy, how does social policy improve welfare in the model, what are the channels by which distortionary effects may undo the benefits of redistribution, and, finally, how will the effectiveness of social policy be measured?

:The redistributive scheme to be evaluated consists of taxation proportional to income, combined with lump-sum transfers. The actual results concerning the effectiveness of the scheme in reducing poverty will not apply directly to all possible social policies, nor even to all those policies that are conceptually similar to that discussed here. However if the results indicate that the equilibrium effects are indeed quantitatively important for the policy in question, then if other policies should prove more effective than the simple scheme considered here, it is likely to be because superior policies have better equilibrium properties. In other words, if the results for the simple policy are significant, this would suggest that the type of analysis developed here is essential for designing and evaluating more effective policies.

3.1 The Benefits of Redistribution

The model as developed above implies two obvious roles for social policy in improving the lot of children born to poor parents. First, the fact that parents are unable to borrow against the future income of their children implies that some poor parents will be forced to choose lower levels of human-capital investment in their children. In the absence of such a constraint, all parents would choose the level of investment that equates the marginal return on human capital to that on bequests. In this case, the only differences in human capital level would arise from differences in the cost of educating children, which in this model is assumed to depend only on the human capital of the parents. The borrowing constraint, by reducing the investment in poor children, raises the marginal return on investment in the human capital of poor children above the equilibrium rate of interest, and therefore redistribution could actually increase per-capita income by re-allocating investment from bequests to the children of the rich to a use with a higher rate of return. More to the point however, redistribution can reduce poverty by increasing the human capital of children from poor families, and as a result, reducing the fraction of adults in the next generation with very low labor income.

The other channel for social policy to reduce poverty is by cushioning the effect of bad luck in the labor market. If parents could, they would insure their children against such bad luck; in the model however, only the rich do this, by bequeathing non-risky financial wealth to their children. The result is that children with poor parents are particularly vulnerable to bad labor-market luck. Social policy can reduce the impact of this outcome by providing transfers.

3.2 The Distortionary Effects of Redistribution

The main drawback to redistribution is the distortionary effects it might have on the decisions of households concerning the quantity and quality of children. Though this has been long recognized, it is only recently that quantitative policy evaluations have begun to incorporate explicit models of the decision problems of the household. In general, redistribution will have distortionary effects through both the taxes required to finance the scheme and through the spending on transfers. In the context of the current model, proportional taxation reduces the return to investing in children, as well as the resources

available for doing so. In addition, proportional taxation reduces the effective wage, and hence the opportunity cost of having children. If parents choose to have more children as a result of the tax, then it is certain that they will invest less per child, and for poor parents, it is likely that this will result in a lower level of human capital per child. The transfers will have similar distortionary effects.

These distortions have some direct effects on poverty, which can be modelled by taking the distribution and prices as given, and allowing parents in the benchmark distribution to choose new decision rules in the light of the social policy. Since parent's human capital is fixed in this type of experiment, the distortionary effects of the policy on income are limited to the effect of fertility on earnings.

The indirect, or long-term effects of these distortions are likely to be much stronger, because parental decisions affect their children's income much more than their own. Hence the model also includes at least two channels by which social policy can have indirect effects on poverty via the distortions outlined above. First, the distributions of households over human capital and wealth depend on the parental fertility and investment decisions. When these change, so do the future distributions; the fraction of people living in poverty may well increase if the result of the policy is to increase the fertility of the poor without increasing the human capital of their children. This type of effect can be modelled by taking the decision rules from the previous exercise and tracking the evolution of the economy over time, under the restriction that the government budget balances in the steady state.

The second indirect channel not only follows from the above shifting of the above distributions, but it also in turn introduces new distortions in parental decisions. This is the response of the equilibrium price level to changes in the capital labor ratio. The capital-labor ratio is determined by the wage and wealth distributions; aggregate capital and labor are equal to the sums over all households of the inherited bequests and effective labor supply, respectively. To model this channel is quite a complex task, because of the interaction between prices and parental decisions on the one hand, and between the distribution and the government deficit on the other. This is the reason for insisting on an equilibrium model in which the return to human capital investment is endogenous. In models in which the return to human capital is independent of aggregate conditions, as is the case when parents are modeled as deriving happiness directly from the education of their children, as in Becker and Tomes (1979), this effect is simply not present.

3.3 Comparing Economies

Although the preceding section outlined the types of economies that should be compared, it is not obvious what the criteria for comparison should be. On the one hand, it would be straightforward to calculate the sort of inequality indices, such as the Gini coefficient, or mobility measures, such as the intergenerational correlation of income, that are standard in the literature. However the object of concern here is not so much inequality or social mobility per se, but the welfare of the poor. Even if the level of per capita income were invariant to policy, inequality would not be a sufficient for comparing poverty, because any given measure of inequality can be consistent with different distributions of households over income, which in turn may imply very different degrees of misery among the poor.

Social policy in the model is likely to result in a trade-off between alleviation of poverty and an increase in its incidence, due to reduced incentives for investment. Welfare in the model depends largely on the level of household income, but parents do not just care about their own income: if a social policy raises the expected utility of the children of poor parents, then this may in itself justify redistribution, even if inequality increases as a result. It is relevant therefore to ask whether poor parents are indeed happier in the regime with redistribution, rather than whether inequality or poverty has increased or decreased.

It is difficult to make welfare comparisons in this environment for a number of reasons. First, since people are not altruistic towards their parents, their actual welfare cannot be justified by reference to their parent's happiness. If the transition path of the economy from one policy regime to the other were to be calculated, then it would be possible to add up the dynastic utility under redistribution of each parent in the original economy. However, it is possible to imagine a policy that made the initial distribution better off in this dynastic sense, but made the long-run descendants much worse.This might be the case for instance if a policy encouraged higher fertility at the expense of less parental investment per child. Since the ancestors are long dead, and their welfare is irrelevant to the descendants, it is difficult to justify the relative misery of the living poor by referring to the benefits enjoyed by their ancestors.

Second, population is endogenous and responds to the social policy regime. Welfare analysis provides no way to evaluate policies or compare economies where the popuation depends on the social policy. A policy that results in

increased fertility and lower consumption per capita, for example, might result in lower average utility in the steady-state but higher total utility (this problem is discussed in Razin and Sadka (1995)). The problem conceptually is that social policies change both the distribution and the number of households in the steady state, but there is no way to aggregate utility across different numbers of households.

The third problem with evaluating policy is that people in the model are heterogeneous along several dimensions; in other words, there are different reasons why people are poor. Some of the poor have high-education but bad market luck, others have low education and average luck; while some have asset income, most do not. One result of these state differences is that the poor make different decisions, especially concerning fertility, and therefore the impact of the policy is likely to vary even among families at the same income percentile.

3.4 A Compensation Measure of Policy Effectiveness

Perhaps the best way to summarize the quandary is to say that it is not clear how the real-world concern with inequality that social policy has been enacted to address translates into a normative ranking of distributions when policy has important equilibrium effects. The approach taken in the analysis below is to describe the steady-state distributions of the benchmark and policy economies in more detail, by comparing the agents at a given income percentile in each economy.

To make this comparison, this paper uses an extension of the concept of compensating variation. This measure is the average amount of one-time additional income transfer required to make an agent at a given percentile rank in the income distribution as well off in the new economy as an agent at the same rank in the benchmark economy. At a given rank in the income distribution however, agents may differ along all three dimensions of the state vector, so the compensating variations are computed for a sample of agents centered on the percentile in question, and the average is reported as a percentage of the average income of parents in the benchmark sample for each percentile rank.

The analysis below reports the distributional outcomes of different policy experiments in terms of these compensations. The compensation is calculated by solving the parental problem of the agent in question subject to a

budget constraint with an extra amount of income[2]. The problem is to find the amount of extra income that gives the agent in the benchmark steady-state economy a maximized value equal to the average value of an agent at the same percentile in the new economy. For each percentile rank x define a neighborhood to be a mass ε of households with median income equal to the income percentile x. Let the average welfare of a household in the neighborhood of percentile x in economy a be represented by $\overline{V}^a(x)$. Consider a household i in the neighborhood of x in economy b. The additional income $\widetilde{y}_i$ required to compensate the agent is defined by:

$$V^b(e_i, a_i + \widetilde{y}_i, z_i) = \overline{V}^a(x) = \frac{1}{\mu^a(\varepsilon(x))} \int_{\varepsilon(x)} V^a(e_j, a_j, z_j)\, d\mu^a(e_j, a_j, z_j)$$

This is a simple fixed-point problem that can be solved by any reliable root-finder algorithm. The benefits of social policy across the income distribution are measured by the average compensation required at a given income percentile:

$$y(x) = \frac{1}{\mu^b(\varepsilon(x))} \int_{\varepsilon(x)} \widetilde{y}_i(e_i, a_i, z_i, x)\, d\mu^b(e_i, a_i, z_i)$$

. Note that this measure has nicer properties than a compensation measure based strictly on consumption for two reasons. First, due to concavity of the utility function, additional income is optimally only partly allocated to additional consumption, so a consumption-based measure overstates the required compensation. Second, due to heterogeneity, people in the neighborhood of a given income percentile may have widely varying consumption levels, and hence the elasticity of their welfare with respect to additional consumption may be equally variable. For both of these reasons, a consumption-based compensation is an unreliable measure of the relative attractiveness of occupying the same income percentile in different economies.

[2]It could be argued that the utility-percentile rank might be more appropriate, but it turns out that income and utility rankings are almost identical. Ranking by labor income on the other hand would drastically change the results, as a significant fraction of the low-earnings population have high non-labor income.

4 Direct Effects of Redistribution

This section presents the results of a series of simulation experiments intended to explore the direct effects, as described above, of a simple social policy. The experiments in this section hold constant the distributions of wealth and earnings, as well as the wage and the interest rate. These are all set identical to the corresponding objects in the benchmark economy. Thus the only effects of social policy on income inequality or poverty are due either to the direct effect of the transfers on the inequality measures, or to the responses of the current generation of parents to the incentive distortions implied by the policy. The analysis decomposes the response by including experiments where parts of the response are suppressed. In the first set of experiments, the fertility decision rule is also fixed as in the benchmark economy, whereas in the second set, fertility is allowed to respond optimally to the new social policy. Each set of experiments includes one where the social policy is temporary, for the current generation only, and one where the policy continues forever. The objective of the exercise is to demonstrate the effectiveness of social policy as it might appear in measurements made in the same generation that the policy is implemented.

The policy consists of a proportional tax on household income, set at 5%, which is used to finance a lump-sum transfer. The balanced-budget requirement implies that the amount of transfer per household will equal 5% of mean household income before redistribution. To the extent that mean income varies according to the experiment, the amount of the lump-sum transfer will also vary by experiment. Only fertility decisions will affect parental income however, so in the fixed-fertility experiments income, and hence the social policy, should be the same.

The analysis here is "short run" in the sense that the wage and the interest rate are fixed as they were in the steady state without policy. Thus is not a true short run of the dynamic model, in the sense that it imposes myopia about price adjustments on the generation that is contemporaneous with the introduction of the policy. However this is close conceptually to the type of effect estimated in empirical analysis that attempts to infer the effects of social programs without fully identifiying the structural parameters of the behavioural equations, as discussed by **?**. Moffit shows that most empirical studies fall into this category of incomplete identification, whether they use time-series or cross-section variation to identify program effects, because of endogeneity of program participation by agents, and because program pa-

rameters themselves are responses to the distribution of households. Thus the contrast in the results between this section and the next may provide guidance for interpretation of these empirical studies.

All of the results are shown in Table 2, which reports the aggregate and summary measures for each economy, and Figures 1-4, which show, respectively, the gains of households, the fertility response, and the education response, all by income percentile, as well as the income distribution.

4.1 Fertility Decisions Fixed

Suppose that the redistribution scheme is temporary: today's parents will be subject to it but the world faced by their children will be the same as the benchmark economy. What will be the effect of the social policy on the households composing the benchmark distribution? The behavioral rules for the descendants will be unaffected by the policy, so parental behavior will only be affected by the direct effects of the tax and the lump-sum subsidy. Taking a 5% tax rate as given, the amount of the redistribution is determined by the condition that the budget balance within the generation that the transfer is affected. This resulted in a transfer of 6.25% of median earnings. The aggregate results are shown in the column for Experiment 1 of Table 1.

The first two rows of the table show the mean tax revenue collected per household, and government spending per household, both expressed as a percentage of the mean household income in the benchmark economy. Mean and median incomes are given in the next two rows, again as a percentage of the mean disposable income in the benchmark economy. Since both the distribution and the fertility decision rules are constant, there is no variation in labor supply across these economies; furthermore, capital holdings are given for the current generation, so neither income nor earnings vary with the social policy.

Since the policy results in net transfers to low-income households, we should see the income Gini decrease in both cases, but no movement in the earnings coefficient. The wealth Gini increases slightly; theoretically, this is due to the small increase in bequests made by the richest households. When the redistribution is permanent, as shown in Experiment 2, the Gini coefficient for wealth increases, because the poorer of the parents who were saving before redistribution no longer require quite as high a level of expected income for their children, and find it more costly to provide income, due to

the tax the children will pay.

Table 1: Direct Effects of Lump-Sum Redistribution

	Bench-	Fertility Fixed		Optimal Fertility	
	Mark	1-time	Permanent	1-time	Permanent
Experiment	0	1	2	3	4
Tax Revenue	0.0%	5.0%	5.0%	5.0%	4.9%
Govt. Spending	0.0%	5.0%	5.0%	5.0%	4.9%
Mean Income	100.0%	100.0%	100.0%	99.6%	97.4%
Median Income	79.2%	80.2%	80.2%	80.2%	78.6%
Earnings Gini	0.37	0.37	0.37	0.37	0.39
Income Gini	0.37	0.35	0.35	0.36	0.37
Wealth Gini	0.73	0.75	0.77	0.75	0.73
Pct. No Bequest	48.5%	53.3%	59.7%	53.9%	48.7%
Pct. Lo-Income	16.4%	13.6%	13.6%	13.6%	20.9%
IG Corr Income	0.54	0.46	0.45	0.52	0.60
Av. Fertility	1.47	1.47	1.47	1.56	1.85
Savings Rate	7.12%	5.80%	5.05%	5.71%	7.22%

The main effects of the policy on the aggregate economy are visible in the following rows, which relate most directly to parental investment. The row marked "Pct. No Beq." shows the percentage of parents who did not leave bequests to their children; a temporary redistribution results in less disposable income among the rich, which induces the poorest of those who made bequests to lower total investment per child[3]. The Gini coefficient for wealth increases more with permanent redistribution because the marginal benefit of investment in children's income falls as the amount of children's income that is independent of parental investment increases. The percentage of parents with less than 1/2 of the benchmark economy's median income are given by the row marked "Pct Lo Inc". Since both temporary and permanent redistributions succeed in reducing this number, the redistribution is effective in reducing poverty by this measure.

[3]Note that this reduction in total investment per child might result in an increase in education investment, because bequests lower the benefit of education.

The following rows show that the redistribution policy succeeds in reducing the intergenerational correlation of income but not that of human capital. In other words, the social policy does not succeed in increasing education investment per child among the poor; any positive welfare effects are due to subsidizing the consumption of the poor, and to insuring everyone's children against bad labor-market outcomes. Since the distribution of children's income does not change with the temporary redistribution, the entire reduction in the correlation must be ascribed to the increased income of poor parents. In other words, it is not the case that a one-time redistribution succeeds in making the children of poor parents better off, merely that the correlation falls because the poor parents no longer appear quite so poor.

After examining these summary measures, the question remains: how much better off are the poor under this social policy, and how much worse off are the rich? To answer this question, it is necessary to turn to a more disaggregate view of the economy, and apply the type of qlefare comparison described in the previous section.

Figure 1 shows the additional income the average household in a neighborhood of percentile x would require in the benchmark economy to attain the average utility of a household at the same percentile in the economy with redistribution. This measure of the welfare gain from redistribution is reported as a percentage of benchmark income, and will be positive when the household is better off under redistribution. Therefore, under a temporary policy, this gain should be positive for all households with less than mean income (before taxes), and negative for households above the mean income.

The line for Experiment 1 corresponds to the direct benefit of a one-time redistribution to the current generation. This line shows that the lowest-income percentile households are much better off under the transfer, as they would require an additional 35% of their benchmark income in order to be indifferent between the two economies. Since all households with less than average income receive a net transfer, all of these households, about 65% of the population, are better off under redistribution. Richer households are worse off, but because they can adjust their spending in response to the tax, they require significantly less than 5% of their original income to compensate them for the redistribution.

The line for Experiment 2 shows that taking into account the effects of redistribution in the future raises the benefits to the very poor significantly higher, to 45% of benchmark income for the very poorest, though the increase is much smaller for households that are less poor. Note that the percentage

of the population that is better off under redistribution has increased to 75%. Except for the very poorest households, the benefits of the redistribution are predominantly due to the increase in the consumption of the parents. For a closer look at this, we turn to Figure 2, which shows the human capital investment policies before and after redistribution, expressed as a percent of the human capital of the children in the benchmark model. The lines for experiments 1 and 2 show little change, except for the poorest 10%, who, when the policy is temporary, choose a human-capital level for their children as much as 15% higher than in the benchmark case. When the redistribution is permanent however, poor parents actually invest less in their children, because the redistribution lowers the expected marginal utility of consumption of the children. The figure also shows an interesting indirect effect: since the rich bequeath less to their children, this raises the marginal benefit of education to their children, which results in higher education of the rich, especially where the parents now choose zero bequest.

The results of these first two experiments show strong direct effects of redistribution on the welfare of the poor households, even though the conventional measures of poverty and inequality respond only slightly. Therefore a disaggregate approach to measuring welfare is essential for evaluating social policy; while redistribution did little to reduce inequality or the incidence of poverty in the model, it was very effective in making the poor better off. Furthermore, the strongest effect of redistribution so far is that it subsidizes adult consumption; there is little gain to the current parents of future redistribution to their children, nor is parental investment in children's human capital significantly increased.

4.2 Fertility Decisions Optimal

Fertility in the preceding experiments was assumed fixed in the sense that parents chose the same number of children they would have chosen without redistribution. In the next two experiments, the parents are allowed to take the social policy into account when choosing fertility. The aggregate results are shown as the columns corresponding to experiments 3 and 4 of Table 1. When the policy is assumed to apply only to the current generation, which is the case in Column 4, there appear three effects of allowing fertility to vary optimally: higher intergenerational correlations, higher fertility, and lower savings rates. The intergenerational correlation of income increases, to 0.52, and that of earnings to 0.45, which is actually higher than in the benchmark

case. Average fertility, which increases to 1.58, is also above the level of the benchmark economy. These effects are magnified when policy is permanent, as in Column 5, but the savings rate, which fell slightly in Experiment 3, actually increases above the benchmark level. The fact that the savings rate increases as a result of the permanent policy indicates that the elasticity of child quantity is greater than that of child quality for the households who are saving for bequests.

Since the fertility choices were non-optimal in the earlier experiments, it must be the case that these changes reflect an even higher benefit of redistribution to some parents. This is readily apparent from Figure 1, by the fact that the compensation line for Experiment 4 lies above that of Experiment 2 for households below the 50th percentile, and coincides thereafter. Figure 2 shows that those households that benefit from being allowed to adjust their fertility choose higher fertility when the policy is continued into the future, whereas if the policy is one-time only, then the line for Experiment 3, shows a much smaller fertility shift, restricted to the households below the 25th percentile.

Two important consequences of fertility adjustment are shown in Figure 3. First, parents who choose higher fertility invariably choose a much lower level of education for their children, This is a result of the budget constraint binding for the households that are so poor as to respond to the social policy through fertility choice. For poor households, the optimal investment per child is quite low, so this makes the fertility margin less costly and hence more responsive to the redistribution than the quality margin. Second, even wealthy parents choose less education for their children, because the parents expect their children to be more fertile if they turn out to have bad luck in the labor market. As a result parents invest less per child, which lowers the cost of having more children. Hence although the direct fertility response to the change in income is quite small, the anticipated effect of future redistribution results in a relatively large fertility increase among the current adults.

Since the distribution of state variables was assumed constant, it is likely that these policy experiments will have little effect on the income distribution. This is confirmed in Figure 4, which shows that when fertility is fixed, redistribution reduces the density to the left of the poverty line, and increases the density around the mean income level. When fertility is optimal however, the density to the left of the poverty line increases, and that to the right of the mean falls, so redistribution results in a fatter left tail, and in

the remainder of the population being more closely packed together around the mean.

To summarize, when fertility is fixed, redistribution is very effective in making the poor better off and imposes low costs on the rich, even though transfers do not increase parental investment, despite the binding borrowing constraint of the poor. However fertility choice in the model tends to undo the inequality-reducing effects of redistribution, even when the distribution of state variables is held fixed. The direct fertility response to increased income is quite small; the main effect of fertility choice on parental investment is due to the anticipated effects of redistribution in the future, which result in both significantly lower investment per child and higher fertility in the present. Hence, the indirect effect of a permanent transfer scheme is a strong reduction in human capital investment by all parents. This accounts for virtually the entire effect of redistribution on children's human capital. This is entirely due to the higher expected fertility of future generations. The savings rate also rises with fertility however, because the elasticity of child quantity is greater than that of investment for the rich.

5 Long-run Effects of Redistribution

The previous experiments took as fixed the distribution of households over the parental state space, as well as the levels of the wage and the interest rate. In this section, these restrictions are relaxed; the experiments that follow explore the effects of social policy on the long-run income distribution. These experiments analyze the long-run distributions that result from lump-sum transfers subject to balancing the government budget, the first with fertility decision-rules fixed, the second with re-optimization of all the decision rules. In order to decompose the interaction between fertility and the social policy, this section also discusses a third experiment, in which the policy consists only of the tax; there is no redistribution. Therefore the analysis in this section permits the comparison of the importance of tax versus transfer effects on the long run outcome, and also the importance of fertility vs price adjustment effects. The results for all of these experiments are shown in Table 2.1 and in Figures 3.5 to 3.8.

5.1 Fertility Decisions Fixed

From Experiments 5 (tax-only) and 6 (tax and transfer), with the fertility decision rules fixed, it is clear that the transfer scheme drastically lowers steady-state income; while mean income under the tax scheme remains close to the benchmark level, adding the transfers to the social policy results in a fall of mean income to 43% of the benchmark level. This is reflected also in the figures for tax revenue, which is 4.5% of mean benchmark income in Experiment 5, but only 2.2% in Experiment 6. Overall, the main effects of the tax policy are to increase both the savings rate and the percentage of parents with very low income by close to 50%. In addition, the intergenerational correlations of earnings and income both fall because parents invest less per child, so parental resources are less important than in the benchmark economy.

Table 2: Equilibrium Effects of Lump-Sum Redistribution

	Bench-	Fertility Fixed		Optimal Fertility	
	Mark	Tax Only	Tax/Tran.	Tax Only	Tax/Tran.
Experiment	0	5	6	7	8
Tax Revenue	0.0%	4.5%	2.2%	3.7%	0.9%
Govt. Spending	0.0%	0.0%	1.9%	0.0%	0.8%
Mean Income	100.0%	84.6%	43.0%	69.8%	18.1%
Median Income	79.2%	67.8%	26.9%	56.1%	10.3%
Earnings Gini	0.37	0.37	0.50	0.36	0.52
Income Gini	0.37	0.37	0.51	0.36	0.50
Wealth Gini	0.73	0.74	0.85	0.73	0.86
Pct. No Bequest	32.2%	52.2%	69.7%	49.5%	77.4%
Pct. Lo-Income	16.4%	24.6%	63.9%	28.9%	87.9%
IG Corr Income	0.54	0.50	0.77	0.53	0.76
Av. Fertility	1.47	1.44	1.82	1.53	2.31
Savings Rate	7.12%	6.66%	4.74%	6.64%	4.62%
(K/L)0	1.95	1.87	1.67	1.75	1.35
(K/L)1	1.95	1.87	1.64	1.70	1.34

The effects of the transfer part of the social policy include, in addition to the drastic income loss referred to above, a larger increase in inequality: the

wealth Gini increases from 0.73 to 0.85, while those for earnings and income increase from 0.37 to 0.5. The intergenerational income correlation increases to 0.77, a 50% increase, because poor parents reduce their investment in kids by more than richer parents do. Average fertility is also significantly higher under the transfer scheme, even though the fertility decision rules are fixed: this is because the policy has increased the density of the population in the high-fertility parts of the distribution over state variables (i.e. the low-education portion of the distribution). This is a fundamental lesson of this approach to analyzing social policy: even if fertility decisions are entirely unresponsive to social policy, the result may be an economy with much higher fertility.

By the standard summary measures therefore, the social policy has failed to reduce poverty: a higher fraction of households live in poverty, and income inequality has increased, as has the intergenerational education correlation. What remains to be seen is the effect of the social policy across the income distribution. Figure 5 shows that the effect of the tax-only policy on welfare is quite uniform across the income distribution: most households in the benchmark economy would be willing to give up between 10 and 20% of their income in order to avoid being transferred to the same percentile in the economy with redistribution. The effect of the transfer policy however is to make the poor much worse off than the rich part of the income distribution, Families below the 70th income percentile would give up 70% of their income to avoid the tax/transfer regime, while the richest 10% of households would only give up slightly more than 40% of their income.

The greater welfare loss of the poor reflects the inequality-increasing property of the transfer scheme alluded to earlier. The mechanism underlying this property is the quantity-quality trade-off. When fertility was higher, investment in education was lower for two reasons: higher parental fertility lowers resources for investment, and higher expected fertility of the children lowers the marginal benefit to the child of an increase in her own human capital. The previous section showed that the latter effect was more important than the former in determining the response of education to social policy. In the current experiments, with the fertility-decision rules fixed, the trade-off still operates, because parents, due to their lower human capital and to the lower capital-labor ratio, face a lower wage, and hence lower time costs of children. They therefore have more children than in the benchmark model, which results in a lower level of investment per child. In Experiment 6, this is true of the poorest half of the population, who can be seen in Figure 7 to choose an

education level 85% to 40% lower than they would have chosen in the benchmark model. Richer parents also choose less education for their children, even though their own fertility is lower than in the benchmark case, because the expected fertility of their children is higher, and hence their expected time on the labor market is lower.

The resulting income distributions are shown in Figure 8. Under the tax-only policy (Experiment 5), the distribution of income is essentially a leftward translation of the benchmark distribution; when transfers are included however (Experiment 6), the distribution flattens considerably,implying higher inequality, in addition to shifting leftward by a much greater amount than in the tax-only case.

5.2 Fertility Decisions Optimal

It is clear from the previous subsection that even when the fertility response to social policy is suppressed, that the long-run effects of the social policy in question are disequalising The remaining question is whether fertility response to the decision rules results in a quantitatively important adjustment to these results. One may argue that since the redistribution has already been shown to fail as a remedy for poverty, that perhaps it unnecessary to run an experiment that is likely only to worsen the apparent effects of the social policy. One reason why such a test is still relevant to evaluating social policy is that one interpretation of the analysis so far is that redistribution has short-run benefits for poor households that are undermined by the long-run disequalizing effects. The cost of the short-run gain therefore depends on the distribution of welfare in the future; if this latter depends on the fertility response, then there may be policies that appear worthwhile only until this response is taken into account.

When fertility responds to the tax-only scheme, as shown in Experiment 7 of Table 2, mean income falls to 69.8% of the benchmark level, and inequality, as measured by the Gini coefficients, barely responds for earnings, income and wealth. Since the tax reduces the return to human capital, two effects are implied: first, the value of parent's labor time is reduced, which increases fertility, and second, the tax implies a lower return to investing in children's human capital; both effects result in lower human capital per child for constrained parents, hence the fall in average income, while only the second effect reduces the human capital of the children of richer parents;

so inequality rises.

Apart from a fall in mean income from 85% to 70% of the benchmark level, the effects of the tax alone (Experiment 7) are relatively insensitive to the way in which fertility is modeled: according to Table 2, allowing fertility to respond actually decreases the percentage of adults who give no bequest, although the percentage of adults with less than half of the benchmark income increases from 24.6% in the fixed-fertility case to 28.5%. Fertility itself increases only slightly on average. However when transfers are added in (Experiment 8), the results are much more dramatic. The mean income falls from 43% of the benchmark level in the fixed-fertility case to 17%. Median income in the new steady state is less than 10% of the mean income in the benchmark, while fertility increases from 1.82 kids per parent in the fixed-fertility case to 2.34. The measures of inequality on the other hand are relatively stable; only the percent of households with low income increases, and that is actually due to the fall in the level of income, since the threshold for low income is defined as 1/2 of median income in the benchmark model.

In Figure 5 it can be seen that the effects on welfare of adding in the fertility response are very large: compensation required to make a household indifferent between the same income percentile ranks increases by 50% for the tax-only case (Experiments 5 and 7), while in the transfer case (Experiments 6 and 8), the increase is about a third for poor households, and it actually doubles for the richest 5% of households. Figure 6 shows that the fertility increase over the fixed-fertility case is concentrated in the poorest 50%, and that education falls much more for the rich than for the poor. The income distribution for Experiment 8, shown in Figure 8, is shifted significantly to the left of that for Experiment 6, and has largely regained the shape of the pre-transfer distribution.

To summarize, the long-run effects of redistribution are even more highly dependent on the fertility response than was the case for the short-run. The mechanism for this dependence appears to be only partly the parental quality-quantity trade-off; parents also invest less in their kids because the kids themselves are expected to be more fertile. Because the incentive effects of redistribution shift the human capital distribution to the left, these interactions between fertility and education are important even when fertility decision rules themselves do not respond to the social policy. Finally the failure of redistribution to reduce poverty is largely due to the transfer scheme itself; taxation to fund the scheme also reduces income and welfare, but the magnitude of this effect is quite minor by comparison with the effect of the

transfers. Furthermore, the taxation scheme does not seem to affect inequality or intergenerational persistence, while these are increased dramatically as soon as the transfer scheme is introduced.

6 Conclusion

The goal of this paper was to assess the importance of dynamic equilibrium effects of social policy on society's ability to alleviate poverty through redistribution. The hypothesis was that such effects would result from the distortionary effects of redistribution on parental decisions concerning fertility and investment per child. Previous literature had suggested that the long-run effects of redistribution could be disequalizing, but the idea of this paper was to ask specifically whether the poor might be better off in the long run, and to arrive at a quantitative comparison of the economies with and without redistribution.

The method chosen to measure these effects was to simulate the effects of redistribution using a dynamic general-equilibrium model of the income distribution, as developed in the previous paper. This model, based on a Becker-Tomes style theory of altruistic parents, generated transmission of income inequality via parental investment in both human capital and bequests, and allowed an important role for luck in the determination of income. The particular parameterization of the model was the same as the benchmark version of the previous paper, which had been chosen so as to match some key features of household income inequality and fertility in U.S. . data. This parametrization was chosen so that the results would reflect the likely importance of the dynamic equilibrium effects in the U.S. .

Social policy was modeled as a proportional income tax combined with lump-sum transfers to households. The main constraint on social policy was that the budget must balance: the average tax revenue must equal the amount of the transfer. Comparisons across economies were obtained by calculating the amount of additional income households in the neighborhood of a given income percentile in the benchmark economy would require to be indifferent to moving to the same percentile in the new economy.

In order to isolate the equilibrium effects of social policy from the direct effects, two sets of experiments were conducted. In the first set, the distribution was taken as given in the benchmark model. In the second set the new steady-state distributions implied by the social policy were found. Each set

of experiments consisted of simulations with fertility decision rules fixed as in the benchmark model and with fertility decision rules responding optimally to the social policy. The first set of experiments, corresponding to the direct or "short-run" effects, was in turn divided into experiments where the policy was applied only to the current generation of adults, and those where the policy persisted into the future. The second set of experiments, where the distribution was allowed to evolve, contrasted the effects of the tax/transfer policy with a social policy consisting only of the tax.

The results of the first set of experiments showed that redistribution was more effective in increasing the welfare of the poor when fertility was flexible, but had strong perverse effects on children's education. All experiments in this set showed that redistribution contributed strongly to the welfare of the poor, who would have required as much as 45% additional income in the benchmark economy in order to be indifferent towards the economy with redistribution. The rich by contrast were only slightly worse off, in the sense that a 2-3% reduction of their income in the benchmark would have made them indifferent between the two economies.

The response of fertility to a one-time redistribution was quite small, but much greater when redistribution was permanent. This suggests that the main effect of redistribution on fertility is via the anticipated effects on children. Even small increases in fertility among the poor resulted in large decreases in their investment in children's human capital. However the anticipated effects of permanent redistribution, namely higher expected fertility of the children, resulted in much lower human capital investment by *all* parents, even the very rich, as the expected marginal benefit of human capital investment is decreasing in children's fertility. An interesting conclusion from this exercise is that fertility has important implications for the effectiveness of social policy, even when there is no fertility response whatsoever to the policy regime.

The second set of experiments showed that redistribution made everyone significantly worse off in the long run, and that this result was magnified when fertility was allowed to respond to the policy. Part of this adverse effect was due to the tax and part was due to the fertility response, but by far the strongest effect was the cumulation over generations of the incentive effect of the transfers themselves on parental investment per child. Poor people in the benchmark economy could have given up 80% of their income, and yet been indifferent towards occupying the same income percentile in the economy with redistribution. The change in the income distribution can be

seen as consisting of a large leftward shift, as income fell to less than 10 % of the average in the benchmark economy, and as an increase in inequality: the Gini coefficients for earnings, income and wealth all increased significantly, the first two by roughly 50% over the benchmark level.

It is clear from the analysis that it can be very misleading to assess the long-run effects of social policies without modelling fertility. The same policies that fail to help the poor in the equilibrium experiment appear much more beneficial to the poor, though still ineffective, under fixed fertility rules. Holding the fertility rules fixed diminishes the apparent cost of redistribution.

Of course the analysis presented here is not the final word on assessing social policy. Even within the context of the current model, it is possible that the short-run benefits of redistribution are more important than the long-run costs, as it may take several generations for the economy to approach the new steady state. In addition, the model was calibrated to match the U.S. . data before the policy experiments were conducted. Therefore the model's results are really about *increases* in social policy. It may be just as realistic to build social policy into the calibration procedure for the benchmark procedure and then compare the steady-state to one with no redistribution. If the model without policy implies that poor families are stuck in a poverty trap, then it is likely that social policy has important benefits even in the long run.

The analysis discussed only one redistributive scheme, so the question is what can one infer from this exercise that may be useful in assessing other, perhaps more realistic social policies? The key point is that equilibrium effects are not only likely to be significant, they are possibly the most important components of the long-run cost of reducing inequality today. Therefore the design and evaluation of any permanent social policies should make heavy use of general-equilibrium modeling: the lack of an empirically significant behavioral response in terms of education or fertility is no guarantee that the policy will not have future indirect effects on these decisions that may be strong enough to wipe out the short-run poverty-reducing benefits.

References

Aiyagari, S., J. Greenwood, and N. Guner (1997): "On the State of the Union," Working paper, University of Rochester.

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 (1979): "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility," *Journal of Political Economy*, 87(6), 1153–1189.

Davies, J. B. (1991): "Does Redistribution Reduce Inequality?," 4(4), 538–59.

Davies, J. B., and P. J. Kuhn (1986): "A Dynamic Model of Redistribution, Inheritance, and Inequality," *Journal of Labor Economics*, 24(2), 324–44.

Gottschalk, P., and T. M. Smeeding (1997): "Cross-National Comparisons of Earnings and Income Inequality," *Journal of Economic Literature*, 35(2), 633–87.

Heckman, J. J., L. Lochner, and C. Taber (1998): "Tax Policy and Human Capital Formation," *American Economic Review*, 88, 293–297.

Huggett, M., and G. Ventura (1997): "On the Distributional Effects of Social Security Reform," University of Western Ontario Department of Economics Research Report 9710.

Knowles, J. (1999): "Can Parental Decisions Explain U.S. Income Inequality?," PhD Dissertation, University of Rochester.

Lucas, R. E. (1976): "Econometric Policy Evaluation: A Critique," *Journal of Monetary Economics*, 2, 19–46.

Moffitt, R. A. (1997): "The Effect of Welfare on Marriage and Fertility: What do We Know and What do we Need to Know?," Institute for Research on Poverty Working Paper No. 1153-97.

Razin, A., and E. Sadka (1995): *Population Economics*. MIT Press, Cambridge, Massachusetts.

Figure 1: Compensation in the Short-Run

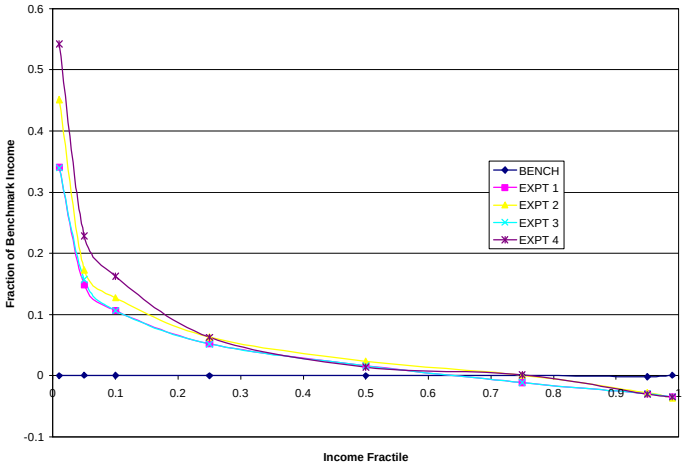

Figure 2: Fertility Response in the Short-Run

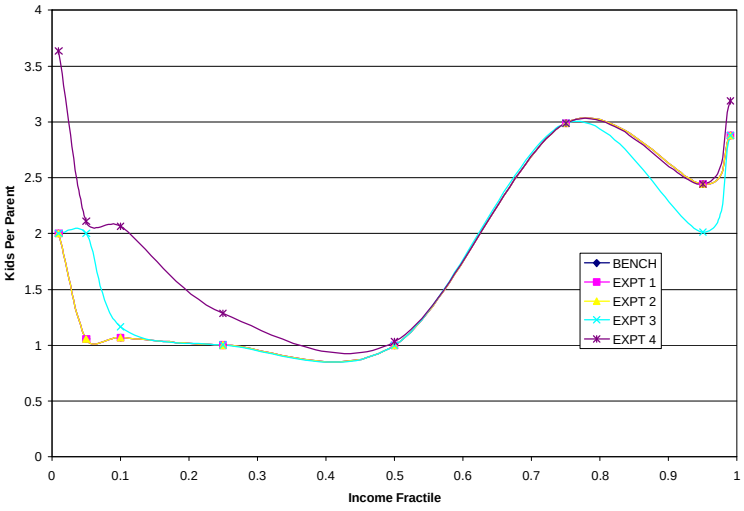

Figure 3: Education Investment in the Short-Run

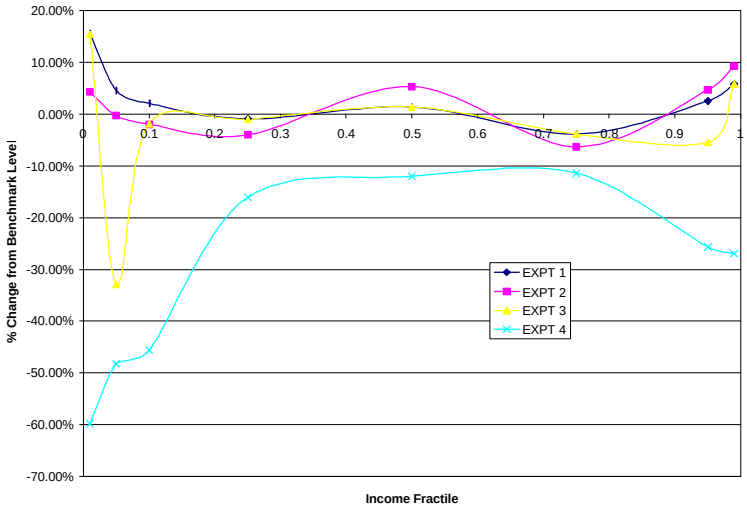

Figure 4: Income Distribution in the Short-Run

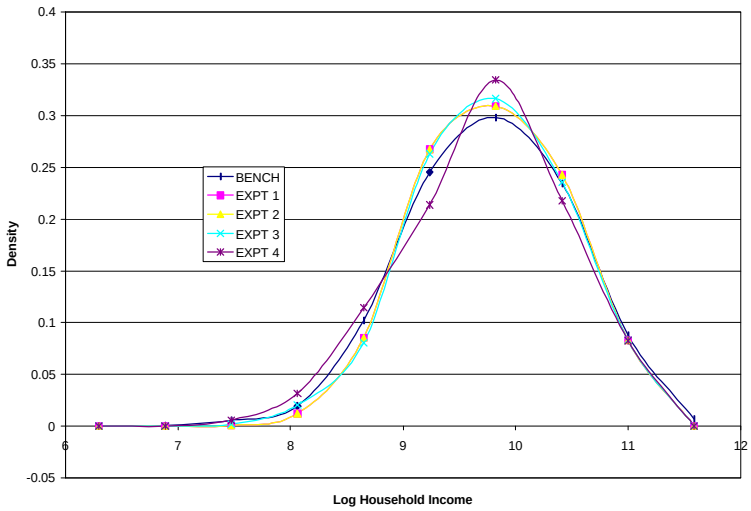

Figure 5: Compensation in the Steady-State

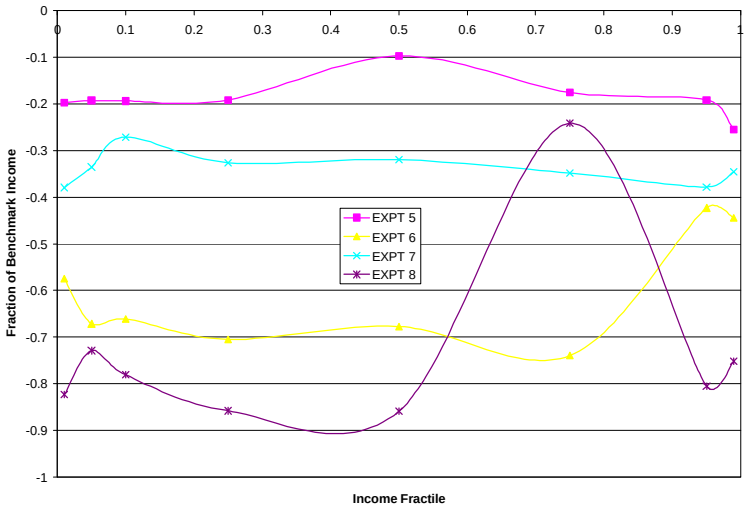

Figure 6: Fertility in the Steady-State

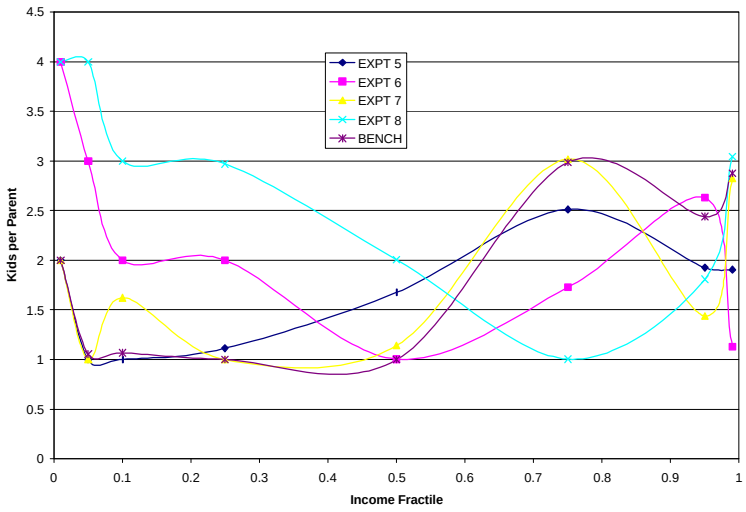

Figure 7: Education Investment in the Steady-State

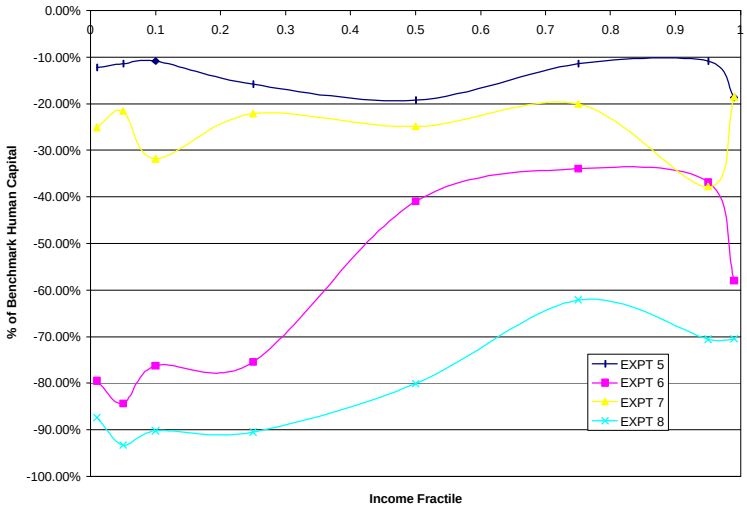

Figure 8: Income Distribution in the Steady-State