NBER WORKING PAPER SERIES

DOES THE CONSUMPTION OF DIFFERENT AGE GROUPS MOVE TOGETHER? A NEW NONPARAMETRIC TEST OF INTERGENERATIONAL ALTRUISM

Andrew Abel

Laurence J. Kotlikoff

Working Paper No. 2490

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 1988

We thank Lawrence Weiss, Robert Barsky, Jinyong Cai and seminar participants at Brown University and at the National Bureau of Economic Research's summer workshop on financial markets for very helpful comments. Jinyong Cai provided outstanding research assistance. This research was funded by the National Science Foundation. Our data are the 1980 through 1984 Consumer Expenditure Surveys which are publicly available through the Bureau of Labor Statistics. The data sets used in the final analysis are available from the authors. The research reported here is part of the NBER's research programs in Financial Markets and Monetary Economics, Economic Fluctuations and Taxation. Any opinions are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #2490
January 1988

Does the Consumption of Different Age Groups Move Together?
A New Nonparametric Test of Intergenerational Altruism

ABSTRACT

In recent years Robert Barro's (1974) ingenious model of intergenerational altruism has taken its place among the major theories of consumption and saving. Despite its policy importance, there have been few direct tests of the Barro model. This paper presents a new direct test that is based on a property of the Barro model that, to our knowledge, has not previously been exploited. This property is that the Euler errors (i.e., disturbances in the Euler equations) of altruistically linked members of Barro extended families (clans) are identical. Under time-separable, homothetic utility, this equality of Euler errors means that, controlling for clan preferences about the age distribution of consumption, the percentage changes over time in consumption of all Barro clan members are equal. With some weak additional assumptions, this proposition implies that the average percentage change in household consumption within an age cohort should be the same for all age cohorts.

Testing the Barro model by comparing average percentage changes in consumption across age cohorts is particularly advantageous because it is nonparametric; in determining whether the average consumptions of different age cohorts move together we place no restrictions on preferences beyond the assumptions of homotheticity and time separability. In particular, each Barro clan can have quite different preference parameters.

The new quarterly Consumer Expenditure Surveys (CES) covering 1980 through the first quarter of 1985 are an excellent data set for determining whether the consumption of different age groups moves together. The CES records the consumption of each sample household for up to four quarters, and thus can be used to determine the average quarterly percentage change in consumption of households in a given age group.

The null hypothesis of our test is that cohort differences in the average percentage change in consumption are due simply to sampling and measurement error. Alternative hypotheses, suggested by the Life Cycle Model, are that (1) the percentage changes in the average consumptions of any two cohorts are more highly correlated the closer in age are the two cohorts, and (2) the variance in the percentage change in consumption is a monotone function of the age of the cohort.

The data fail to reject the null hypothesis of equal Euler errors. Indeed, the results provide fairly strong support for the intergenerational altruism model as opposed to the Life Cycle Model.

Andrew B. Abel
Department of Finance
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104

Laurence J. Kotlikoff
NBER
1050 Massachusetts Avenue
Cambridge, MA 02138

In recent years Robert Barro's (1974) ingenious model of intergenerational altruism has taken its place among the major theories of consumption and saving. The model, which starts with the simple assumption that parents care about the welfare of their children, yields the remarkably strong conclusion that, apart from distorting marginal incentives, deficits and all other government redistributions between generations have no effect on the economy. The possibility that deficits, unfunded social security, and similar policies do not matter has received considerable attention.

Despite its policy importance, there have been few direct tests of the Barro model. The main difficulty in directly testing the model at the micro level is the lack of data detailing both the consumption and resources of altruistically linked households. Direct tests of the model with macro data are also problematic because they require the aggregation of different Barro clans (sets of altruistically linked households) each of which may have a different utility function.

This paper presents a new direct test of the Barro model. The test is based on a property of the Barro model that, to our knowledge, has not previously been exploited. This property is that the Euler errors (i.e., disturbances in the Euler equations) of altruistically linked members of Barro extended families (clans) are identical. Assuming utility is homothetic and time separable, this equality of Euler errors means that, controlling for clan preferences about the age distribution of consumption, the percentage changes over time in consumption of all Barro clan members are equal. Intuitively, since consumption of each clan member is based on overall clan resources, and not the distribution of resources over clan members, any shocks to the resources of specific clan members will be spread across all clan members. Under the homotheticity and time separability assumptions, spreading shocks over all clan members means changing the consumption of all members by the same percentage.

Ideally, one would test this proposition by simply comparing changes in the consumption of different clan members. Unfortunately, the requisite clan-specific data is not available; indeed, it may be very difficult to determine who is and who is not a member of a particular altruistically-linked clan. As indicated by Kotlikoff (1983) and Bernheim and Bagwell (1985), clans may be quite large because of current as well as potential future intermarriage.

Let us admit, however, the possibility of multiple clans, but assume for the moment that each clan has the same age structure. In this case, the average Euler error within each age cohort will be equal, since the error of each clan will receive the same weight in each of the cohort averages. However, the assumption of identical age structures within each clan seems too strong. A weaker assumption that leads to the same result is a zero correlation between the age structure of clans and their Euler error; i.e., the fact that a clan accounts for a larger than average fraction of households in an age group does not help predict how its Euler error will differ, on average, from the average clan Euler error.

Testing the Barro model by comparing average cohort percentage changes in consumption is particularly advantageous because it is nonparametric; in determining whether the average consumption of different age cohorts moves together we place no restrictions on preferences beyond the assumptions of homotheticity and time separability. In particular, each Barro clan can have quite different preferences.

The new quarterly Consumer Expenditure Surveys (CES), which, to date, are available from the first quarter of 1980 through the first quarter of 1985, are an excellent data set for determining whether the consumption of different age groups moves together. The CES records the consumption of each sample household for up to four quarters, and thus can be used to determine the average quarterly percentage change in consumption of households in a given age group.

The null hypothesis of our test is that cohort differences in the average percentage change in consumption are due simply to sampling and measurement error. Alternative hypotheses, suggested by the Life Cycle Model, are that (1) the percentage changes in average consumptions of any two cohorts are more highly correlated the closer in age are the two cohorts, and (2) the variance in the percentage change in consumption is a monotone function of the age of the cohort.

The data fail to reject the null hypothesis of equal Euler errors in favor of the alternative hypotheses for our definition of total consumption. Indeed, the results provide fairly strong support for the intergenerational altruism model as opposed to the Life Cycle Model.

The paper proceeds in the next section by reviewing briefly the empirical literature bearing on the Barro hypothesis. Section III presents the Barro model and develops the proposition that Euler errors are equal for all clan members. Section IV derives a statistical model to test this proposition. Section V describes the data. Section VI contains the empirical results, and Section VII concludes the paper with some suggestions for additional research.

Section II. Empirical Research Bearing on the Barro Hypothesis

The largest body of empirical literature bearing on the Barro hypothesis relates the consumption time series to the time series of unfunded social security. Chief among these studies are those of Feldstein (1974), Barro (1977), Darby (1977), and Leimer and Lesnoy (1981). Studies relating the consumption time series to other aspects of fiscal policy include Feldstein (1982), Kormendi (1983), and Aschauer (1985). The results of this body of research can be summarized with one word, ambiguous. Even were the results all in agreement, it would be difficult to know precisely what had been learned; as pointed out by Auerbach and Kotlikoff (1983) and Williamson and Jones (1983), if the Life Cycle model is taken as the null hypothesis in these studies, the models are misspecified because of the inability to aggregate the behavior of different age groups. The Auerbach-Kotlikoff paper shows that the regression procedures would reject the Life Cycle Model even using data taken from a pure Life Cycle economy. An alternative view of these regressions is that the Barro model is the null hypothesis. But in this case the regressions also seem to be misspecified both because of aggregation and because they ignore the government's intertemporal budget constraint.

A different body of literature that is relevant to the Barro model as well as other neoclassical models are the Euler equation studies of Hall (1978), Hall and Mishkin (1982), Flavin (1981), Shapiro (1984), Zeldes (1985), Mankiw, Rotemberg, and Summers (1982), Lawrence (1983), Altonji and Siow (1987) and others. These papers test intertemporal expected utility maximization, specifically its implication that the Euler error is uncorrelated with previous information. A rejection of this null hypothesis would rule out the Barro model as well as other neoclassical consumption models. But as stressed by King (1983), tests of the Euler equation require specifying the explicit form of preferences, and rejection of the Euler equation may simply reflect an incorrect choice of preferences. The time series tests of the Euler equation provide mixed results.

While providing evidence that a minority of households are liquidity constrained, most micro level studies appear to accept the Euler equation restriction for the majority of households. For example, both Zeldes (1985) and Lawrence (1983) use the limited consumption data in the PSID and reach the conclusion that the Euler equation holds for the great majority of households.

A recent paper by Boskin and Kotlikoff (1986) directly tests the implication of the Barro model that the age-distribution of resources doesn't affect the age distribution of consumption. They reject the proposition that aggregate consumption is invariant to their proxy for the age distribution of resources. The Boskin-Kotlikoff results should not, however, be viewed as definitive; their analysis, like other time series studies, is subject to aggregation bias. In addition, their specification of preferences and uncertainty, specifically the bivariate distribution of interest rates and labor earnings, may be inappropriate.

Section III. The Equal Euler Error Proposition

Let $U_{i,k,t}$ stand for the utility of household k at time t in Barro clan i. The combined assumptions of homotheticity and time separability imply that utility is of the isoelastic form. Hence, we write $U_{i,k,t}$ as:

(1) $$U_{i,k,t} = \sum_{a=0}^{D} P_{i,k,t,a}\ \theta_{i,k,a}\ \frac{c_{i,k,t,a}^{1-\gamma_i}}{1-\gamma_i}$$

where $P_{i,k,t,a}$ is the number of members age a in household k, clan i at time t, D is the maximum age of life, $\theta_{i,k,a}$ is the weight household k in clan i places on the utility of members age a, and $c_{i,k,t,a}$ is the consumption of the members of clan i who are in household k and are age a at time t. Let $C_{i,k,t}$ stand for total consumption of household k in clan i at time t. Then (1) will be maximized subject to:

(2) $$\sum_{a=0}^{D} P_{i,k,t,a}\ c_{i,k,t,a} = C_{i,k,t}$$

This implies:

(3) $$U_{i,k,t} = \frac{[\sum_{a=0}^{D} P_{i,k,t,a}\ \theta_{i,k,a}^{1/\gamma_i}]^{\gamma_i}\ C_{i,k,t}^{1-\gamma_i}}{1-\gamma_i} = \frac{\beta_{i,k,t} C_{i,k,t}^{1-\gamma_i}}{1-\gamma_i}$$

The Barro clan i's infinite horizon expected utility can now be written as:

(4) $$V_{i,t} = E_t \sum_{s=t}^{\infty} \alpha_i^{s-t} \sum_{h=1}^{N_{i,s}} \beta_{i,h,s} \frac{C_{i,h,s}^{1-\gamma_i}}{1-\gamma_i}$$

where $N_{i,s}$ is the number of households in clan i at time s.

The Barro clan chief maximizes (4) subject to:

(5) $$W_{i,t+1} = W_{i,t}(1+r_{i,t}) + e_{i,t} - C^*_{i,t}$$

where, $C^*_{i,t} = \sum_{h=1}^{N_{i,t}} C_{i,h,t}$, is total clan i consumption at time t. The term $e_{i,t}$ stands for the possibly uncertain labor earnings of the clan at time t; $r_{i,t}$ is the possibly uncertain rate of return earned by clan i at time t, and $W_{i,t}$ is clan i's net worth at time t. In addition to $e_{i,t}$ and $r_{i,t}$, $\beta_{i,h,s}$ for s > t in (4) may be uncertain at time t due to life span uncertainty and uncertainty about clan fertility. Maximization of (4) subject to (5) implies the static conditions:

(6) $$\beta_{i,k,t}\, C_{i,k,t}^{-\gamma_i} = \beta_{i,h,t}\, C_{i,h,t}^{-\gamma_i}$$

and the intertemporal conditions:

(7) $$\alpha_i \beta_{i,k,t+1} C_{i,k,t+1}^{-\gamma_i}(1+r_{i,t+1}) = \beta_{i,k,t}\, C_{i,k,t}^{-\gamma_i}\epsilon_{i,k,t+1} ,$$

where $\epsilon_{i,k,t+1}$ is the Euler error with a mean of unity. Equations (6) and (7) imply that the Euler errors of all Barro households within a clan are identical; i.e.:

(8) $$\epsilon_{i,k,t+1} = \epsilon_{i,h,t+1} = \epsilon_{i,t+1}$$

Section IV. <u>A Test of the Equal Euler Error Proposition Based on Cohort Data</u>

In this section we develop a method of testing the equal Euler error proposition using cohort data. We start by taking logarithms of (7), yielding:

(10)

$$\log(C_{i,k,t+1}/C_{i,k,t}) = (1/\gamma_i)\log(\alpha_i\beta_{i,k,t+1}/\beta_{i,k,t}) - \log(\epsilon_{i,t+1}/(1+r_{i,t+1}))$$

Consider all households in clan i whose heads are age a. Take the average of (10) over all such households. The resulting average of equation (10) is given by equation (11) where we define the averages of the left hand side and the two terms on the right hand side of (10) respectively by:

$$(11) \quad Y^a_{i,t+1} = \psi^a_{i,t+1} + \mu_{i,t+1}$$

Note that the term $\mu_{i,t+1}$ is not indexed by age since the Euler errors of each of clan i's households are identical. Next average (11) over all clans. This produces (12) where $s^a_{i,t}$ is the fraction of age a households that belong to clan i at time t, and M is the total number of clans.

$$(12) \quad \sum_{i=1}^{M} s^a_{i,t} Y^a_{i,t} = \sum_{i=1}^{M} s^a_{i,t} \psi^a_{i,t} + \sum_{i=1}^{M} \mu_{i,t} /M + \sum_{i=1}^{M} s^a_{i,t}(\mu_{i,t} - \sum_{j=1}^{M} \mu_{j,t} /M)$$

In (12) the cohort average value of $\mu_{i,t}$ is written as the simple unweighted average of the Euler errors across all clans (the second term on the right hand side of the equation) plus the cohort average value (weighted by each clan's fraction of all cohort households) of the deviation of the clan's Euler error from the unweighted average Euler error over all clans. We make the assumption that this third term on the right hand side, which is the population covariance between a clan's Euler error and its share of its population in the age group, is zero.

We can rewrite the remaining terms in (12) more compactly by letting $\bar{Y}^a{}_t$ denote the left hand side of (12), $\bar{\psi}^a{}_t$ denote the first term on the right hand side of (12), and $\bar{\mu}_t$ denote the second term on the right hand side.

$$(12') \quad \bar{Y}_t{}^a = \bar{\psi}_t{}^a + \bar{\mu}_t$$

Equation (12') states that the cohort average value of the percentage change in consumption (more precisely, the log of the ratio of consumption at t+1 to consumption at time t) equals a term, $\bar{\psi}^a{}_t$, which depends on age and time, plus a term $\bar{\mu}_t$, which is independent of age.

Because of sampling and measurement error, the true population mean, $\bar{Y}_t{}^a$, is not perfectly observable. Hence, in (13), we set the observed population-weighted sample mean of the logarithm of the ratio of consumption at time t+1 to consumption at time t, $\hat{\bar{Y}}_{ta}$, equal to the true population mean, $\bar{Y}_t{}^a$, plus a term, $\eta_t{}^a$ that reflects sampling and measurement error. Our null hypothesis is that $\eta_t{}^a = \omega_{at} / h_t{}^a$, where ω_{at} is an independently and normally distributed random variable with mean zero and variance σ^2, and $h^a{}_t$ adjusts for the sampling error in our weighted estimate of $\bar{Y}_t{}^a$. Specifically, $h_t{}^a$ equals $\Sigma_k w_{atk}{}^2 / (\Sigma_k w_{atk})^2$, where w_{atk} is the CES population weight at time t for household k in cohort a. In (12') the term $\bar{\psi}^a{}_t$ reflects the average growth in consumption due to demographic changes in household composition.

Since we are dealing with data over only a five year interval, in (13) we drop the time subscript and treat $\bar{\psi}^a{}_t$ as a time invariant, but age-specific constant.

$$\hat{\bar{Y}}^a_t = \bar{\psi}^a + \bar{\mu}_t + \eta^a_t \tag{13}$$

Equation (13) forms the basis for our statistical test of the equality of average cohort percentage changes in consumption. Under the null hypothesis of equal Euler errors ω_{at} is i.i.d. across ages a and time periods t with variance equal to σ^2. If the null hypothesis fails to hold and the weighted average Euler errors differ across age cohorts, the error term $\eta^a{}_t$ will capture not only measurement and sampling noise, but also each cohort's time t average Euler error after controlling for age and time effects. Our alternative hypothesis is, therefore, that ω_{at} is not simply i.i.d., but depends on age as specified below:

$$\begin{aligned} &E(\omega_{it}\omega_{js}) = 0 \text{ if } s \neq t \\ &E(\omega_{it}\omega_{jt}) = \rho^{|i-j|}\, \sigma^2\, \gamma^{i+j} \end{aligned} \tag{14}$$

According to (14) the variance of ω_{it} increases or decreases with age depending on whether γ exceeds or falls short of unity, and the correlation of ω_{it} and ω_{jt} for $i \neq j$ depends on the the size of the age gap, $|j-i|$. For example, if ρ exceeds zero, (14) says that the correlation of ω_{it} and ω_{jt} for age groups i and j is larger the closer in age are the age groups i and j. The case in which $\rho=0$ and $\gamma=1$ corresponds to the null hypothesis. Values of ρ

and γ as well as the age and time effects in (13) are estimated by maximum likelihood. The Appendix presents the likelihood function under the null and alternative hypotheses.

To see how the alternative hypothesis might hold in a nonaltruistic life cycle model, consider the case of an individual who (1) lives to time T, (2) has a logarithmic utility function, (3) receives an income stream in each period that evolves as a random walk, (4) faces a zero rate of interest, and (5) has a zero rate of time preference. Under these assumptions the individual's time t+1 Euler error, ϵ_{t+1}, can be related to the error in the random walk process, ϕ_{t+1}, by considering the individual's lifetime budget constraint. Since the present value of realized lifetime consumption must equal the present value of realized lifetime resources, it is also true that the expected present value of lifetime consumption equals the expected present value of lifetime resources. In addition, the difference in the expected present value of lifetime consumption at times t+1 and t equals the difference in the expected present value of lifetime resources at times t+1 and t. This last relationship and the successive application of the Euler equation under logarithmic utility, namely that $C_{t+1} = C_t/\epsilon_{t+1}$ (where C_t is the individual's time t consumption), yields the following relationship between ϵ_{t+1} and ϕ_{t+1}:

(15) $$C_t(E_{t+1} - E_t) \sum_{i=1}^{T-t} \prod_{s=1}^{i} \frac{1}{\epsilon_{t+s}} = (T-t)\, \phi_{t+1} \qquad \text{or}$$

(15′) $$C_t \sum_{i=1}^{T-t} \left(\frac{1}{\epsilon_{t+1}} E_{t+1} \prod_{s=2}^{i} \frac{1}{\epsilon_{t+s}} - E_t \prod_{s=2}^{i} \frac{1}{\epsilon_{t+s}} \right) = (T-t)\, \phi_{t+1}$$

In considering either (15) or (15') it is important to note that while the expectation of ϵ_{t+1} conditional on information at time t equals unity, the time t expectation of $1/\epsilon_{t+1}$ exceeds unity, by Jensen's inequality. The relationship between ϵ_{t+1} and ϕ_{t+1} involves the expectation of the product of the reciprocals of future Euler errors. Since the reciprocals of the Euler errors are not necessarily independently distributed, these expectations do not lead to a simple relationship between ϵ_{t+1} and ϕ_{t+1}. In addition, since these expectations are conditional on the age of the individual, T-t, the variance of ϵ_{t+1} can increase or decrease with age even for this case in which the future income process is independent of age.

Of course, there is no reason to believe that the shocks to full future resources (including those arising from random rates of return) have the same variance independent of age. Hence, even ignoring the complex expectations of equations (15) and (15'), there is, in the context of the Life Cycle Model, no particular reason to believe that the variance of Euler errors will be independent of age. While the variance of Euler errors may rise or fall with age under the Life Cycle Model, one would expect in that model that the variance of individuals close in age would be quite similar; i.e., one would expect a gradual change by age in the variance of the Euler error. In addition, one would expect a positive correlation in Euler errors for individuals close in age because they will experience similar shocks. The alternative assumption specified in (14) permits both a correlation in Euler errors between age groups that is larger the closer together are the age groups as well as a variance in Euler errors that gradually rises or falls with age.

Section V. The Data

The ongoing Consumer Expenditure Survey (CES), which began in the first quarter of 1980, interviews approximately 4500 households in each quarter. Most households are interviewed four times in the Consumer Expenditure Survey. The four interviews always ask a common set of questions about consumption, but some other questions are asked only in the first and fourth interviews, and others are asked only in the fourth interview. Some households are interviewed fewer than four times because they drop out of the sample. Others are interviewed fewer than four times because of the sample design; in an effort to maintain in each quarter the same fraction of households responding to a first, second, third, and fourth interview, the CES administers the second, third, or fourth interviews to some households as their initial interview. If the household's initial interview is a second interview, the household will be interviewed two more times. If a household's initial interview is a third interview, the household will be interviewed once more. And if the household's initial interview is a fourth interview, the household will not be reinterviewed.

The approximately 4500 interviews in each quarter are spread over each month of the quarter. In the interviews households are asked about their consumption expenditures in the previous three months. Hence, a household interviewed in January of 1981 reports consumption expenditures for October, November, and December of 1980, while a household interviewed in March of 1981 reports consumption expenditures for December of 1980 and January and February of 1981. Unfortunately, for most expenditure items, households only report total expenditures in the previous three full months and do not provide a month-to-month breakdown of those expenditures. As a consequence, the data for a household interviewed, say in January, cannot readily be combined with data from a household interviewed in February since the two quarterly observations cover overlapping, rather than identical quarters. In effect, each wave of the Consumer Expenditure Survey provides three overlapping sets of observations on quarterly consumption. In our analysis we treat each of the three quarterly data sets separately and refer to them as quarterly samples 1, 2, and 3.[1] For purposes of analyzing the quarterly data we considered 58 age cohorts corresponding to ages 23 through 80.

Given the lumpiness of some nondurable consumption expenditures, such as vacation trips, it is useful to test the equal Euler error proposition with semi-annual as well as quarterly data. For those households who were interviewed four times, the four quarterly observations can be combined to form observations on semi-annual consumption. There are 6 possible semi-annual data sets. For example, households interviewed in January, April, July, and October in year t provide an observation on the ratio of consumption over the period April-September in year t to consumption over the period October in year t-1-March in year t. Households interviewed in July and October of year t and January and April of year t+1 provide an observation on the ratio of consumption over the period October in year t-March in year t+1 to consumption over the period April-September in year t. These types of observations produce a single data set of semi-annual changes in consumption. One can also form a data set using households interviewed for the first time of four times in April and other households interviewed for the first of four times in October. Hence, the April–July–October–January sequence provides two semi–annual data sets. The May–August–November–February sequence provides another two semi–annual data sets; and the June–September–December–March sequence provides the final two semi–annual data sets.

Because of the smaller number of households who completed all four surveys, we constructed three–year age cohorts; i.e., we combined ages 23, 24, and 25 into one age group, ages 26, 27, and 28 into another age group, etc. up to the age group covering ages 77, 78, and 79. This difference in the definition of an age cohort should be kept in mind in comparing the quarterly and semi–annual results presented in the next section; because of the difference in definitions one would expect the estimated values of ρ and γ based on the semi–annual data to be roughly the cube of their respective values based on the quarterly data.

The definition of aggregate consumption used in this study is total consumption expenditures excluding expenditures on housing, insurance, and consumer durables. We exclude housing both because adjustments to housing consumption are infrequent and because it is very difficult to impute quarterly or semi–annual rent accurately for homeowners. Insurance expenditures were excluded because such expenditures represent risk pooling as opposed to consumption per se. In addition, the data records both negative and positive amounts of insurance expenditures, where a negative amount corresponds to a claim payment. Expenditures on durables should clearly be excluded from the definition of consumption. In contrast, imputed rent should be included; unfortunately, data on the stocks of durables are not sufficient for that purpose.

The CES provides population weights in each quarter for each household interviewed. These weights depend on the age of the household head as well as other economic and demographic characteristics. We use the time t+1 sample weights in determining the cohort-specific weighted average value of the logarithm of the ratio of consumption at time t+1 to consumption at time t; i.e., we construct a weighted value of $\overline{Y}_{at}$.

Households that reported less than $150 of quarterly expenditure on food were excluded from the sample. This is the only form of sample selection in our analysis. Some limited analysis indicated that including households with very small quarterly food expenditure would not materially alter the results.

Section VI. <u>Empirical Findings</u>

As a prelude to examining estimates of ρ and γ, Figure 1 considers how the age-consumption profile changed over the period 1980 through 1984. Ignoring demographic change, the proposition that each cohort's consumption should change, on average, by the same percentage, implies a constant age-consumption profile. In forming Figure 1 we calculated the annual weighted average of quarterly consumption (measured in 1985 dollars) at each individual age for households interviewed in April, August, and December of each of the five years. We combined these weighted averages within each calender year to produce annual values of average consumption by age of the household. Next we divided annual consumption in year t at each age by the average consumption of 45 year old households in year t. Finally, we smoothed these relative consumption values for each year by regressing them against an intercept and a fourth order polynomial in age. In these regressions the R^2 values each exceed .9. Figure 1 plots the resulting five smoothed polynomials of consumption at a particular age relative to consumption at age 45.

The curve with the most dashes corresponds to 1980, the curve with the second most dashes corresponds to 1981, etc. The curves in the Figure suggest that the age-consumption profile flattened out in 1983 and 1984. Compared with 1981, for example, the 1984 relative consumption of 60 year olds is over 10 percent larger. The $F(20,264)$ value for the test that the five polynomials are the same is 17.94, greatly in excess of the 5 percent critical value of 1.66. Since the changes in the shape of the relative age-consumption profile do not appear to be due to changes in demographics, it provides some evidence against the Barro Model; however, unlike the next set of findings, these profiles consider levels, not changes in consumption, and, as such, do not control as well for the composition of the sample; i.e., the levels of consumption of the elderly in 1983 and 1984 may reflect samples whose older households happened to belong to clans with greater total resources.

Table 1 also provides some preliminary data analysis of the null hypothesis. This Table compares quarterly changes in consumption of different age groups for quarterly sample 1. The corresponding tables for quarterly samples 2 and 3 are quite similar. For purposes of Table 1 we consider five broad age categories: 23-29, 30-39, 40-49, 50-59, 60-69, and 70 +. For each of these six age groups we report quarterly values of 100 times the deviation of the weighted average $\hat{\bar{Y}}_t^a$ from the mean value, $\sum_{t=1}^{19} \hat{\bar{Y}}_t^a$, taken over the 19 quarters in our sample. According to (12′) and ignoring measurement and sampling error, these deviations, which we refer to as average adjusted Euler errors, should be identical for each of the six age groups.

Table 1 indicates that these average adjusted Euler errors are typically very different across the six age categories. There are only 5 of 19 quarters in which the signs of the adjusted Euler errors are the same for each age group. For quarterly samples 2 and 3, there are only six of 19 and five of 19 such quarters, respectively. Even in quarters when all the adjusted Euler errors have the same sign, there is still a considerable difference in the magnitude of the errors. For example, in the third quarter of 1980 all Euler errors are negative, but the error for the age group 23–29 is less than a fifth the size of the error for the age group 50–59. Such large differences between the smallest and largest average Euler error arise in each of the quarters for each of the three quarterly samples.

Another informal way to assess the data is to regress $\hat{\bar{Y}}_t^a$ on a set of age group dummies and time dummies, either quarterly or semi-annual. The results from this regression can be compared with the results from regressing the same dependent variable on age dummies and the interaction of each of the time dummies with each of the age dummies.[2] According to (13), given a particular time period t, the age-time interactions should have identical coefficients. For purposes of this regression using quarterly data, we constructed six age dummies corresponding to the six age groups of Table 1. The F values for quarterly samples 1, 2, and 3 are 1.470, 1.237, .746, respectively. Since the $F(90,987)$ 5 percent critical value is 1.27, the age-time interactions are significantly different in only one of the three quarterly samples. The $F(25,103)$ values in the corresponding regressions for the six semi-annual samples are .912, 2.414, 2.538, 1.485, 1.768, and .612. The 5 percent critical value in this case is 1.61. Hence, age-time interactions are significantly different in only three of the six semi-annual samples.

Table 2 presents our maximum likelihood estimates for ρ and γ for the three quarterly data sets based on individual age cohorts from age 23 through age 80. None of the reported estimates of these parameters is significantly different from the values predicted by the null hypothesis of intergenerational altruism. Indeed, in the case of γ, two of the three estimates are equal to 1 to four decimal places and the third value of 1.0020 implies that the variance of ω_{it} for 70 year olds is only about 15 percent larger than the corresponding variance for 20 year olds. One of the three point estimates for ρ is negative; a negative value of ρ, even were it significant, seems highly unlikely from the perspective of the Life Cycle Model. The other two nonnegative values of ρ suggest a very small correlation between the consumption of adjacent age groups. Even if these estimates were significant, their values seem quite small.

The likelihood functions associated with Table 2 are rather sharply peaked; hence, one can reject values of ρ and γ that are substantially different (in an economic sense) from the maximum likelihood estimates. Table 3 presents the range of values of ρ and γ that fall within 95 percent Chi-squared confidence intervals around the maximum likelihood estimates.[3] According to the Table, even if one takes the largest values of ρ and γ that cannot be rejected by the data, the resulting estimates provide no strong evidence of substantial departure from the null hypothesis of intergenerational altruism.

Since many of the consumption expenditures included in our definition of nondurable expenditure may not be made each quarter, the results in Table 2 may, in part, reflect the lumpiness of nondurable expenditures; i.e., the variance in consumption changes due to the lumpiness of expenditures may dominate the results. Hence, it may be useful to repeat the analysis using simply food expenditures which is much less lumpy than, for example, clothing expenditures or vacation trips. The results based on quarterly food expenditures are quite similar to the results based on total nondurable-nonhousing consumption expenditure. The point estimates in the three samples of ρ are -.0590, -.0680, and .0020. The point estimates of γ in the three samples are 1.0010, .9980, and 1.0030. The estimates of ρ and γ are not jointly significantly different from 0 and 1, respectively; the respective χ^2 values for the three samples for the joint test that ρ equals 0 and γ equals 1 are 4.087, 4.991, and 4.757 - all of which lie below the 5 percent critical value of 5.991.

Another way to consider the lumpiness of expenditures is to repeat the analysis with semi-annual data. Table 4 presents the results based on the six semi-annual consumption data sets, which, as mentioned, combine three ages into a single age cohort. Once again, none of the estimates of ρ and γ are separately or jointly significantly different from the null hypothesis values of $\rho = 0$ and $\gamma = 1$. Three of the six point estimates of γ lie above 1 and three lie below 1. Three of the six point estimates of ρ are positive and three are negative. Hence, like the quarterly estimates, there is no suggestion in the data that the null hypothesis is strongly disfavored. Unlike the quarterly results, however, several of the estimates of γ are economically more important. For example, the estimate for γ in the sixth sample of 1.0220 implies that the variance of ω_{it} for very old households is over 1.7 times the variance for very young households. In addition, for each of the six samples the confidence intervals around γ include economically significant as well as economically insignificant values. Thus the semi-annual results do not provide as strong evidence against the Life Cycle Model as do the quarterly results.

One might question whether we have properly controlled for demographic change in treating $\bar{\psi}^a$ as an age-specific time-invariant constant. One way to consider whether the results are sensitive to treatment of demographics is to re-estimate the model defining household consumption as household consumption per household member or per adult equivalent in the household; in forming adult equivalents we treat each child under age 18 as equal to .5 adults. We tried each of these alternative definitions of household consumption. The quarterly results are essentially the same as those in Table 2.[4] The semi-annual results are only slightly different from those in Table 4; when consumption is measured either as household consumption per member or per equivalent adult, the null hypothesis is rejected in only two of the six semi-annual samples.[5]

Section VII. <u>Conclusion</u>

After controlling for demographic change and sampling and measurement noise, the average change in consumption appears to be identical across all age groups. This rather strong finding is suggested by Barro's model of intergenerational altruism, in contrast to the Life Cycle Model. An important attribute of these results is that they are nonparametric in nature; specifically, in comparing the average change in consumption across age groups we place no restrictions on preferences beyond the assumptions of homotheticity and time separability.

It may be that quarterly changes and even semi-annual changes in consumption reflect quite lumpy expenditures and that testing the equal Euler error proposition on annual or even biannual would be more appropriate. Unfortunately, appropriate data for such an analysis do not currently exist.

Tabel 1. Quarterly Estimates of Average Adjusted Euler Errors

Quarterly Sample 1

Age Group	1980.3	1980.4	1981.1	1981.2	1981.3	1981.4
23-29	-.226	.551	.744	-.962	-.593	.084
30-39	-.694	.451	.501	-.741	.569	.195
40-49	-.810	1.076	-.008	-.928	-.616	.498
50-59	-1.259	.273	.146	.230	-.495	.101
60-69	-.873	.416	.602	.340	-.390	-.586
70+	-1.05	.396	1.245	-.391	-2.003	1.333

Age Group	1982.1	1982.2	1982.3	1982.4	1983.1	1983.2
23-29	-.895	.569	.088	.187	.395	-.385
30-39	-.229	-.405	-.320	-.042	1.072	-.446
40-49	-.386	-.311	-.078	.153	.817	-.760
50-59	-.295	-.027	-.425	-.202	.799	-.146
60-69	.147	-1.583	-.012	-.423	1.075	-.908
70+	-1.317	2.208	-1.567	-.004	1.910	-.943

Age Group	1983.3	1983.4	1984.1	1984.2	1984.3	1984.4
23-29	.380	.369	-.152	-.741	-.097	.699
30-39	-.200	.073	.200	-.636	.190	.090
40-49	-.132	.859	.249	-.569	.271	.271
50-59	.683	.442	.187	.044	.365	.184
60-69	.269	.664	.058	.231	.391	-.130
70+	-.735	.052	-.012	.880	-.145	-.034

Age Group	1985.1
23-29	-.011
30-39	.373
40-49	.405
50-59	-.604
60-69	.710
70+	.176

Table 2.

Maximum Likelihood Estimates and x^2 Values - Quarterly Samples

	Consumption Measured Per Household						
	Unconstrained		ρ=0		γ=1		ρ=0 γ=1
Sample	ρ	γ	γ	x^2	ρ	x^2	x^2
Sample 1	.0200	1.0020	1.0020	.958	.0031	1.307	2.345
Sample 2	-.0240	1.0000	1.0000	.611	-.0240	.000	.611
Sample 3	.0310	1.0000	1.0000	1.066	.0310	.000	1.066

5 percent critical values for x^2 are 5.991 for two restrictions and 3.841 for one restriction.

Table 3.

95% Confidence Intervals for ρ and γ - Quarterly Estimates

	Consumption Measured Per Household					
	Maximum Liklihood		γ Range		ρ Range	
Sample	ρ	γ	High	Low	High	Low
Sample 1	.0200	1.0020	1.0050	.9990	.0950	-.0300
Sample 2	-.0240	1.0000	1.0030	.9970	.0200	-.0990
Sample 3	.0310	1.0000	1.0030	.9970	.1060	-.0440

Table 4.

Maximum Likelihood Estimates and x^2 Values - Semi-annual Samples

Consumption Measured Per Household

Consumption Category	Unconstrained		$\rho=0$		$\gamma=1$		$\rho=0\ \gamma=1$
	ρ	γ	γ	x^2	ρ	x^2	x^2
Sample 1	-.1330	1.0070	1.0090	2.531	-.1360	.409	3.101
Sample 2	-.1420	1.0080	1.0090	2.945	-.1420	.478	3.490
Sample 3	.0870	.9900	.9910	.024	.0800	.706	1.585
Sample 4	.0550	.9910	.9920	.467	.0490	.555	.918
Sample 5	-.0720	.9840	.9840	.719	-.0690	2.377	3.054
Sample 6	.0500	1.0220	1.0220	.353	.0510	3.679	4.048

5 percent critical values for x^2 are 5.991 for two restrictions and 3.841 for one restriction.

Notes

1. Quarterly sample 1 corresponds to households interviewed in April, July, October, and January. Quarterly sample 2 corresponds to households interviewed in May, August, November, and February. Quarterly sample 3 corresponds to households interviewed in June, September, December, and March. In constructing the data for quarterly sample 1, as an example, we form ratios of a) the July reported quarterly consumption to the April reported quarterly consumption, b) the October reported quarterly consumption to the July reported quarterly consumption, c) the January reported quarterly consumption to the October reported quarterly consumption, and d) the April reported quarterly consumption to the January reported quarterly consumption. In forming the average logarithm of the ratio of consumption say in January 1983 to consumption in October 1982, all households who were surveyed in both October 1982 and January 1983 were included.

2. To illustrate these two regressions, recall that we consider six age groups and there are 19 time periods. Then the initial regression is:

$\hat{\bar{Y}}_t^a = \sum_{i=1}^{6} \delta_i A_i + \sum_{j=1}^{19} \tau_j T_j$, where the δ's and τ are coefficients, A_i is a dummy for age group i, and T_j is a dummy for time period j. The alternative model is: $\hat{\bar{Y}}_t^a = \sum_{i=2}^{6} \delta_i A_i + \sum_{j=2}^{19} \sum_{i=1}^{6} \lambda_{ij} T_j$, and the test is that $\lambda_{ij} = \lambda_{kj}$ for all i, j, and k.

3. These bounds were constructed by holding one of these parameters fixed at its maximum likelihood value and varying the other parameter until the resulting likelihood was significantly (at the five percent level) different from the maximum likelihood.

4. With household consumption defined as consumption per person the point estimates for ρ for quarterly samples 1, 2, and 3 are .0340, -.0430, and .0440, respectively. For γ the corresponding point estimates are 1.0010, 1.0000, and 1.0000. With household consumption defined as consumption per equivalent adult, the three point estimates for ρ are .0320, -.0350, and .0390, while the three point estimates for γ are 1.0010, 1.0000, and 1.0000. The x^2 values for testing the null hypothesis that ρ equals 0 and γ equals 1, are 2.094, 1.952, and .395 for the three quarterly samples when consumption is measured per person, and 2.240, 1.352, and 1.667 for the three quarterly samples when consumption is measured per equivalent adult.

5. With household consumption defined as consumption per person the point estimates for ρ for semi-annual samples 1, 2, 3, 4, 5, and 6 are -.1260, .0450, -.0840, -.1360, .0570, and -.0060, respectively. For γ the

corresponding point estimates are 1.0010, .9710, .9700, 1.0140, .9970, and 1.023. With household consumption defined as consumption per equivalent adult, the six point estimates for ρ are -.1270, .0770, -.0900, -.1310, .0480, and .0100, while the six point estimates for γ are 1.0090, .9730, .9740, 1.015, .9940, and 1.0210. The six respective semi-annual χ^2 values for testing the null hypothesis that ρ equals 0 and γ equals 1 are 3.0763, 7.3924, 8.710, 3.807, .493, and 4.220 when consumption is measured per person, and 3.0441, 6.699, 7.284, 3.735, .533, and 3.581 when consumption is measured per equivalent adult.

References

Altonji, J. and A. Siow, "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data," National Bureau of Economic Research Working Paper, no. 2012, September 1986.

Ando, Albert and Modigliani, Franco. "The 'Life Cycle' Hypothesis of Saving: Aggregate Implications and Tests," American Economic Review, March 1963, 53(1, Part 1), pp. 55-84.

Aschauer, David A., "Fiscal Policy and Aggregate Demand," American Economic Review, vol. 75, no. 1, March 1985, 117-27.

Auerbach, Alan J. and Laurence J. Kotlikoff, "An Examination of Empirical Tests of Social Security Savings," in Social policy evaluation: An economic perspective. Eds: Elhanan Helpman et al. NY: Academic Press, 1983.

Barro, Robert J. "Are Government Bonds Net Wealth?" Journal of Political Econ., Nov./Dec. 1974, 82(6), pp. 1095-1117.

Bernheim, Douglas and Kyle Bagwell, "Is Everything Neutral?", mimeo 1985.

Boskin, Michael J. and Laurence J. Kotlikoff, "Public Debt and U.S. Saving: A New Test of the Neutrality Hypothesis", Carnegie-Rochester Conference Volume Series, February, 1986.

Campell, John Y. and Angus S. Deaton, "Is Consumption Too Smooth," National Bureau of Economic Research Working Paper no. 2134, January 1987.

Darby, Michael R. Effects of Social Security on income and the capital stock. Wash., DC: Amer. Ent. Inst., 1979.

Feldstein, Martin S., "Social Security, Induced Retirement, and Aggregate Capital Accumulation," Journal of Political Economy, Sept./Oct. 1974b, 82(5), pp. 905-26.

Feldstein, Martin S., "Government Deficits and Aggregate Demand," Journal of Monetary Economics, vol. 9, January 1982, 1-20.

Flavin, Margorie A. "The Adjustment of Consumption to Changing Expectations about Future Income," Journal of Political Economy, October 1981, 89(5), pp. 974-1009.

Hall, Robert E. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy, December 1978, 86(6), pp 971-87.

Hall, Robert E. and Frederic Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," Econometrica, vol. 50, 461-81.

Hayashi, Fumio. "The Effects of Liquidity Constraints on Consumption: A Cross Sectional Analysis." Cambridge, MA: NBER Working Paper No. 882, April 1982.

King, Mervin, "The Economics of Savings," NBER working paper, 1983.

Kotlikoff, Laurence J., "Altruistic Extended Family Linkages, A Note", mimeo, National Bureau of Economic Research, 1983.

Kormendi, Roger, "Government Debt, Government Spending, and Private Sector Behavior," American Economic Review, vol. 73, no. 4, December 1983, 994-1010.

Lawrence, Emily, "Do Transfers to the Poor Reduce Savings," mimeo 1983.

Leimer, Dean R. and Selig D. Lesnoy, "Social Security and Private Saving: A Reexamination of the Time Series Evidence Using Alternative Social Security Wealth Variables," Working Paper no. 19, Social Security Administration, Office of Research and Statistics, November 1980.

Mankiw, N. Gregory; Rotemberg, Julio J. and Summers, Lawrence H. "Intertemporal Substitution in Macroeconomics." Quarterly Journal of Economics, 1983.

Miron, Jeffrey A., "Seasonal Fluctuations and the Life Cycle - Permanent Income Model of Consumption," Journal of Political Economy, vol. 94, no. 6, December 1986, 1258-80.

Shapiro, Matthew, "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data," Economics Letters, vol. 14, 1984, 93-100.

West, Kenneth D., "The Insensitivity of Consumption to News About Income," National Bureau of Economic Research Working Paper no. 2252, May 1987.

Williamson, Samuel H. and Jones, Warren L. "Computing the Impact of Social Security Using the Life Cycle Consumption Function." American Economic Review, December 1983, 73(5), pp. 1036-52.

Zeldes, Stephen, "Consumption and Liquidity Constraints: An Empirical Investigation," Rodney L. White Center for Financial Research, Wharton School, University of Pennsylvania, Working Paper no. 24-85, November 1985.

Appendix

The Likelihood Function and the Derivation of Parameter Estimates

Under the assumption that the η_t^as are normal and independent across time, the log of the likelihood function, L, is given by:

$$\text{(A1)} \qquad L = \log K - \frac{1}{2} \sum_{t=1}^{T} \log|V_t| - \frac{1}{2} \sum_{t=1}^{T} \eta_t' V_t^{-1} \eta_t$$

where $\eta_t = Y_t - \psi - \mu_t i$. The term Y_t is an N by 1 column vector, where N is the the number of age cohorts (58 in the case of quarterly data and 19 in the case of semi-annual data) in our data, whose elements are $\hat{\bar{Y}}_t^a$. The vector ψ captures the time-invariant, age-specific constants arising in equation (13) when $\bar{\psi}_t^a$ is time-invariant. The vector i is a column vector of 1s. The term T equals the number of time periods in our data set.

The matrix V_t equals $H_t' V H_t$, and V is defined by:

$$V = \sigma^2 \begin{vmatrix} \gamma^2 & \rho\gamma^3 & \cdot & \cdot & \cdot & & \\ \rho\gamma^3 & \gamma^4 & & \cdot & & & \\ \cdot & & & & \cdot & & \\ \rho^{N-2}\gamma^N & \rho^{N-3}\gamma^{N+1} & & & \cdot & \gamma^{2N-2} & \rho\gamma^{2N-1} \\ \rho^{N-1}\gamma^{N+1} & \rho^{N-2}\gamma^{N+2} & & & & \rho\gamma^{2N-1} & \gamma^{2N} \end{vmatrix}$$

and

$$H_t = \begin{vmatrix} h_{1t} & 0 & \cdots & 0 \\ 0 & \cdot & & \vdots \\ \vdots & & \cdot & 0 \\ 0 & \cdots & 0 & h_{nt} \end{vmatrix} \qquad \text{where } h_{at} \text{ equals } \sum_k w_{atk}^2 \Big/ \Big(\sum_k w_{atk}\Big)^2 \text{ and}$$

where w_{atk} is the CES population weight of household k which is age a at time t.

The first-order conditions resulting from maximizing (A1) with respect to ψ, μ_t, and σ^2 are given respectively in (A2), (A3), (A4):

(A2) $$\sum_{t=1}^{T} V_t^{-1} \eta_t = 0$$

(A3) $$i'V_t^{-1}\eta_t = 0$$

(A4) $$NT = \sum_{t=1}^{T} \eta_t' V_t^{-1} \eta_t$$

From (A3) we have:

(A5) $$\mu_t = (i'V_t^{-1}i)^{-1} i'V_t^{-1}(Y_t - \psi)$$

Normalizing the sum $\sum_{t=1}^{T} \mu_t$ to zero yields:

(A6) $$\sum_{t=1}^{T} (i'V_t^{-1}i)^{-1} i'V_t^{-1}(Y_t - \psi) = 0$$

Equations (A5) and (A2) imply:

(A7) $$\sum_{t=1}^{T} V_t^{-1}(I - (i'V_t^{-1} i)ii')(Y_t - \psi) = 0$$

Multiplying (A6) by i and adding the resulting expression to (A7) leads to:

(A8)

$$\hat{\psi} = [\sum_{t=1}^{T} (V_t^{-1} - (i'V_t^{-1}i)^{-1}(V_t^{-1} - I)ii'V_t^{-1}]^{-1}[\sum_{t=1}^{T} (V_t^{-1} - (i'V_t^{-1}i)^{-1}(V_t^{-1} - I)ii'V_t^{-1})Y_t$$

Given knowledge of the V_ts, we can use (A8) plus (A5) to determine estimates of the μ_ts and the elements of ψ. Rather than solve analytically for the estimates of γ and ρ, we searched over a grid of alternative pairs of these parameters. For each choice of these parameters we formed the V_t matrices and used (A8) and (A5) to calculate the corresponding values of ψ and the μ_ts.