Volume Title: The Distributional Aspects of Social Security and Social Security Reform

Volume Author/Editor: Martin Feldstein and Jeffrey B. Liebman, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-24106-8

Volume URL: http://www.nber.org/books/feld02-1

Conference Date: October 21-23, 1999

Publication Date: January 2002

Title: Social Security and Inequality over the Life Cycle

Author: Angus S. Deaton, Pierre-Olivier Gourinchas, Christina Paxson

URL: http://www.nber.org/chapters/c9750

4

Social Security and Inequality over the Life Cycle

Angus Deaton, Pierre-Olivier Gourinchas, and Christina Paxson

4.1 Introduction

This chapter explores the consequences of Social Security reform for the inequality of consumption across individuals. The basic idea is that inequality is at least in part the consequence of individual risk in earnings or asset returns. In each period, each person gets a different draw, of earnings or of asset returns, so that whenever differences cumulate over time, the members of any group will draw further apart from one another, and inequality will grow. Inequality at a moment of time is the fossilized record of the history of personal differences in risky outcomes. Any institution that shares risk across individuals, the U.S. Social Security system being the case in point, will moderate the transmission of individual risk into inequality, and it is this process that we study in this chapter. Note that we are not concerned here with what has been one of the central issues in Social Security reform, the distribution *between* different generations over the transition. Instead, we are concerned with the equilibrium effects of

Angus Deaton is the Dwight D. Eisenhower Professor of International Affairs and professor of economics and international affairs at the Woodrow Wilson School of Public and International Affairs, Princeton University, and a research associate of the National Bureau of Economic Research. Pierre-Olivier Gourinchas is assistant professor of economics at Princeton University and a faculty research fellow of the National Bureau of Economic Research. Christina Paxson is professor of economics and public affairs at the Woodrow Wilson School of Public and International Affairs, Princeton University, and a research associate of the National Bureau of Economic Research.

Angus Deaton and Christina Paxson acknowledge support from the National Institute on Aging through a grant to the National Bureau of Economic Research, and from the John D. and Catherine T. MacArthur Foundation within their network on inequality and poverty in broader perspectives. All the authors are grateful to Martin Feldstein, Laurence J. Kotlikoff, Jeffrey B. Liebman, and James M. Poterba for helpful comments and discussions.

different Social Security arrangements on inequality among members of any given generation.

A concrete and readily analyzed example occurs when the economy is composed of autarkic permanent income consumers, each of whom has an uncertain flow of earnings. Each agent's consumption follows a martingale (i.e., consumption today equals expected consumption tomorrow), and is therefore the cumulated sum of martingale differences, so that if shocks to earnings are independent over agents, consumption inequality grows with time for any group with fixed membership. The same is true of asset and income inequality, although not necessarily of earnings inequality: see Deaton and Paxson (1994), who also document the actual growth of income, earnings, and consumption inequality over the life cycle in the United States and elsewhere. An insurance arrangement that taxes earnings and redistributes the proceeds equally, either in the present or the future, reduces the rate at which consumption inequality evolves. With complete insurance, marginal utilities of different agents move in lockstep, and consumption inequality remains constant. Social Security pools risks and thus limits the growth of life-cycle inequality. Reducing the share of income that is pooled through the Social Security system, as envisaged by some reform proposals (such as the establishment of individual accounts with different portfolios or different management costs) but not by others (such as setting up a provident fund with a common portfolio and common management costs) increases the rate at which consumption and income inequality evolve over life in a world of permanent-income consumers. Even if inequality is not inherited from one generation to the next, and each generation starts afresh, partial privatization of Social Security will increase average inequality. While much of the discussion about limiting portfolio choice in new Social Security arrangements has (rightly) focused on limiting risk, such restrictions will also have effects on inequality.

Provided the reform is structured so that the poor are made no worse off, it can be argued that the increase in inequality is of no concern (see, e.g., Feldstein 1998), so that our analysis would be of purely academic interest. Nevertheless, the fact remains that many people—perhaps mistaking inequality for poverty—find inequality objectionable, so that it is as well to be aware of the fact if it is the case that an increase in inequality is likely to be an outcome of Social Security reform. There are also instrumental reasons for being concerned about inequality; both theoretical and empirical studies implicate inequality in other socially undesirable outcomes, such as low investment in public goods, lower economic growth, and even poor health (Wilkinson 1996).

The paper is organized as follows. Section 4.2 works entirely within the framework of the permanent income hypothesis (PIH). We derive the formulas that govern the spread of consumption inequality, and show how inequality is modified by the introduction of a stylized Social Security scheme. The baseline analysis and preliminary results come from Deaton and Paxson (1994), which should be consulted for more details, refinements, and reservations, as well as for documentation that consumption inequality grows over the life cycle, not only in the United States, but at much the same rate in Britain and Taiwan. The PIH is convenient because it permits closed-form solutions that show explicitly how Social Security is related to inequality. However, it is not a very realistic model of actual consumption in the United States, and it embodies assumptions that are far from obviously appropriate for Social Security analysis—for example, that consumers have unlimited access to credit, and that intertemporal transfers leaving the present value of lifetime resources unchanged have no effect on consumption. In consequence, in section 4.3, we consider richer models of consumption and saving that incorporate both precautionary motives for saving and borrowing restrictions. These models help replicate what we see in the data, which is the tendency of consumers to switch endogenously from buffer-stock behavior early in life to life-cycle saving behavior in middle age. The presence of the precautionary motive and the borrowing constraints breaks the link between consumption and the present value of lifetime resources, which both complicates and enriches the analysis of Social Security. Legal restrictions prevent the use of Social Security as a collateral for loans, and for at least some people such restrictions are likely to be binding.

Solutions to models with precautionary motives and borrowing constraints are used to document how Social Security systems with differing degrees of risk sharing affect inequality. We first consider the case in which all consumers receive the same rate of return on their assets. Our results indicate that systems in which there is less sharing of earnings risk—such as systems of individual accounts—produce higher consumption inequality both before and after retirement. An important related issue is whether differences across consumers in rates of return will contribute to even greater inequality. Somewhat surprisingly, we find that allowing for fairly substantial differences in rates of return across consumers has only modest additional effects on inequality. The bulk of saving, in the form of both Social Security and non–Social Security assets, is done late enough in life so that differences in rates of return do not contribute much to consumption inequality.

4.2 Social Security and Inequality under the Permanent Income Hypothesis

Section 4.2.1 introduces the notation and basic algebra of the permanent income hypothesis, while section 4.2.2 reproduces from Deaton and Paxson (1994) the basic result on the spread of consumption and income inequality over the life cycle. Both subsections are preliminary to the main analysis. Section 4.2.3 introduces a simplified Social Security system in an infinite horizon model with PIH consumers and shows how a Social Security tax at rate τ reduces the rate of spread of consumption inequality by the factor $(1 - \tau)^2$. Section 4.2.4 discusses what happens when there is a maximum to the Social Security tax, and section 4.2.5 extends the model to deal with finite lives and retirement and shows that the basic result is unaffected.

4.2.1 Preliminaries: Notation and the Permanent Income Hypothesis

It is useful to begin with the algebra of the PIH; the notation is taken from Deaton (1992). Real earnings at time t are denoted y_t. Individual consumption is c_t and assets A_t; when it is necessary to do so we shall introduce an i suffix to denote individuals. There is a constant real rate of interest r. These magnitudes are linked by the accumulation identity

$$A_t = (1 + r)(A_{t-1} + y_{t-1} - c_{t-1}) \ . \tag{1}$$

Under certainty equivalence, with rate of time preference equal to r, and an infinite horizon, consumption satisfies the PIH rule, and is equal to the return on the discounted present value of earnings and assets:

$$c_t = \frac{r}{1 + r} A_t + \frac{r}{1 + r} \sum_{k=0}^{\infty} \frac{1}{(1 + r)^k} E_t(y_{t+k}) \tag{2}$$

for expectation operator E_t, conditional on information available at time t. It is convenient to begin with the infinite horizon case; we deal with the finite horizon case in section 4.2.5.

That consumption follows a martingale is made evident by manipulation of equation (2):

$$\Delta c_t = \eta_t \equiv \frac{r}{1 + r} \sum_{k=0}^{\infty} \frac{1}{(1 + r)^k} (E_t - E_{t-1}) y_{t+k} \ . \tag{3}$$

"Disposable" income y_t^d is defined as earnings plus income from capital:

$$y_t^d = \frac{r}{1 + r} A_t + y_t \ . \tag{4}$$

Saving is the difference between disposable income and consumption,

$$s_t = y_t^d - c_t \ , \tag{5}$$

which enables us to rewrite the PIH rule in equation (2) in the equivalent form (see Campbell 1987):

$$s_t = -\sum_{k=1}^{\infty} \frac{1}{(1 + r)^k} E_t \Delta y_{t+k} \tag{6}$$

Assets are linked to saving through the identity (implied by equations [1] and [5]):

$$\Delta A_t = (1 + r)s_{t-1} \tag{7}$$

Finally, it is convenient to specify a stochastic process for earnings, y_t. It is convenient to do this by assuming that

$$\alpha(L)(y_t - \mu) = \beta(L)\varepsilon_t \tag{8}$$

for lag operator L and polynomials $\alpha(L)$ and $\beta(L)$ and white noise ε_t. As written, and under the usual conditions on the roots, earnings is stationary (around μ) and invertible. In fact, we can allow a unit root in $\alpha(L)$ with essentially no modification. (In the more realistic models in section 4.3, we will work with a process with a unit root but specified in logarithms.)

Given equation (8), we can derive explicit forms for the innovation to consumption (Flavin 1981):

$$\Delta c_t = \eta_t = \frac{r}{1 + r} \cdot \frac{\beta\left(\frac{1}{1 + r}\right)}{\alpha\left(\frac{1}{1 + r}\right)} \varepsilon_t , \tag{9}$$

so that consumption is a random walk and the innovation variance of consumption is tied to the innovation variance of earnings by the autocorrelation properties of the latter.

4.2.2 Spreading Inequality

Begin with the simplest illustrative case, in which earnings are white noise, and add an i suffix for an individual

$$y_{it} = \mu_i + \varepsilon_{it} = \mu_i + w_t + z_{it} , \tag{10}$$

where μ_i is the individual-specific mean of earnings, w_t is a common (macro) component, and z_{it} is an idiosyncratic component. The macro component w_t is also independently and identically distributed (i.i.d.) over time. Given equation (10), equation (3) implies

$$c_{it} = c_{it-1} + \frac{r}{1 + r}(w_t + z_{it}) . \tag{11}$$

As a result, if the idiosyncratic components are orthogonal to lagged consumption in the cross-section (which need not be true in each year but is true on the average by the martingale property), the cross-sectional variance of consumption satisfies

(12)
$$\mathrm{var}_t(c) = \mathrm{var}_{t-1}(c) + \frac{\sigma_z^2 r^2}{(1+r)^2}$$
$$= \mathrm{var}_0(c) + \frac{t\sigma_z^2 r^2}{(1+r)^2}$$

so that consumption inequality is increasing over time.

Note that although equation (12) is derived for the variance of consumption, the increase in consumption variance is general, not specific to a particular measure of inequality. According to equation (11), the household distribution of consumption at t is the distribution of consumption at $t - 1$ plus uncorrelated white noise. Given that the mean is not changing, the addition of noise implies that the distribution of consumption at t is second-order stochastically dominated by the distribution of consumption at $t - 1$, so that any transfer-respecting measure of inequality, such as the Gini coefficient, the Theil inequality measure, or the coefficient of variation (but not necessarily the variance in logarithms), will show an increase of inequality over time.

In the i.i.d. case, saving is given by (see equation [6])

(13)
$$s_{it} = \frac{\varepsilon_{it}}{1+r},$$

while assets satisfy

(14)
$$A_{it} = A_{it-1} + \varepsilon_{it-1}.$$

Because disposable income is the sum of consumption and saving, the change in disposable income satisfies

(15)
$$\Delta y_{it}^d = \frac{r\varepsilon_{it}}{1+r} + \frac{\varepsilon_{it}}{1+r} - \frac{\varepsilon_{it-1}}{1+r} = \varepsilon_{it} - \frac{\varepsilon_{it-1}}{1+r},$$

which implies, after some manipulation, that

(16)
$$\mathrm{var}_t(y^d) = \mathrm{var}_{t-1}(y^d) + \frac{\sigma_z^2 r^2}{(1+r)^2}$$
$$= \mathrm{var}_0(y^d) + \frac{t\sigma_z^2 r^2}{(1+r)^2}.$$

Because the consumption variance is spreading, and because saving is stationary by equation (13), disposable income variance must spread at the same rate as the consumption variance. Note that earnings variance is constant given the stationarity assumption in equation (10), so that

(17)
$$\mathrm{var}_t y = \sigma_\mu^2 + \sigma_z^2 = \text{constant}.$$

From equation (14), the variance of assets satisfies

$$\text{var}_t(A) = \text{var}_0(A) + t\sigma_z^2 . \tag{18}$$

The rate of spread of the variance of assets is the variance of the idiosyncratic component of the innovation of earnings. At a real interest rate of 5 percent, this is 400 times faster than the rate of spread of the variance of consumption and of disposable income. From any given starting point, asset inequality among a group of individuals grows much faster than does consumption or disposable income inequality.

In the United States, the data on consumption, earnings, and income are consistent with the predictions of the theory. Deaton and Paxson (1994) use repeated cross-sections from the Consumer Expenditure Survey (CEX) to trace birth cohorts through the successive surveys, and find that cross-sectional consumption inequality for any given birth cohort increases with the age of the cohort. For example, the Gini coefficient for family consumption (family income) increases (on average over all cohorts) from 0.28 (0.42) at age twenty-five to about 0.38 (0.62) at age fifty-five. We shall return to these findings in section 4.3.

4.2.3 Social Security and the Spread of Inequality

Suppose that the government enacts a simple Social Security system. A proportionate tax on earnings is levied at rate τ, and the revenues are divided equally and given to everyone. We think about the (partial) reversal of this process as a stylized version of reform proposals that pays some part of each individual's Social Security tax into personal saving accounts; the precise mechanisms will be presented in section 4.3.2. We recognize that the establishment of personal accounts has other effects, some of which are not captured under our simple assumptions. Our concern here, however, is with the reduction in the pooling or risk sharing that is implied by removing a part of Social Security tax proceeds from the common pool and placing it in individual accounts. Such accounts provide smoothing benefits for autarkic agents who would not or cannot save on their own account, but they reduce the risk-sharing elements of the current system unless they are supplemented by other specifically risk-sharing features such as transfers from successful to unsuccessful investors.

Because of the infinite horizon and certainty equivalence assumptions, dividing up the revenues and returning them immediately is the same as giving them back later. The model assumes no deadweight loss. Denote before-tax earnings as y_{it}^b and retain the notation y_{it} for after-tax income, $(1 - \tau)y_{it}^b$. In the i.i.d. case we have

$$y_{it} = (1 - \tau)(\mu_i + \varepsilon_{it}) + \tau\bar{\mu} , \tag{19}$$

where the last term is the average revenue of the tax, which is given back to everyone. Equation (19) can also be written as

(20) $$y_{it} = \mu_i - \tau(\mu_i - \bar{\mu}) + (1 - \tau)\varepsilon_{it} .$$

Compared with the original earnings process in equation (10), there is a shift toward the grand mean—the redistributional effect of the Social Security system—together with a scaling of the innovation by $1 - \tau$, which is the risk-sharing component of the Social Security system. The redistribution will change consumption levels for everyone not at the mean, but will not affect the innovation of consumption equation (11), saving equation (13), asset equation (14), or disposable income equation (15), *except* that the original innovation must be rescaled by $1 - \tau$. In consequence, the variances of consumption, disposable income, and assets all evolve as before, but at a rate that is $(1 - \tau)^2$ times the original rate. If the Social Security tax is 12.4 percent, inequality (measured by the variance) will spread at 76.7 percent of the rate that it would in the absence of the system. Imagine an economy in equilibrium, with no inheritance of inequality and no growth in lifetime resources, so that the cross-sectional profile of consumption by age is identical to the lifetime profile of consumption for each cohort, and all consumption inequality is within-cohort inequality. With a working life of forty years and consumption variance originally growing at 5 percent, the imposition of a Social Security tax at 12.4 percent will reduce the cross-sectional standard deviation of consumption by a factor of 5.

In equations (19) and (20), we have not explicitly distinguished the macro common component of the innovation w_t from the idiosyncratic component ε_{it}. If we substitute to make the decomposition explicit, equation (20) becomes

(21) $$y_{it} = \mu_i - \tau(\mu_i - \bar{\mu}) + (1 - \tau)z_{it} + w_t ,$$

which shows that the common component is not insured. The change in consumption warranted by equation (21) is

(22) $$\Delta c_{it} = \frac{r}{1 + r}\Big[w_t + (1 - \tau)z_{it}\Big],$$

but only the second term in the brackets contributes to the spread in consumption variance, and the results are as stated previously.

4.2.4 Social Security with a Maximum

The PIH is not well suited to modeling a Social Security system in which taxes are paid only up to the Social Security maximum. The nonlinearity complicates the forecasting equations for earnings and eliminates the analytical tractability that is the main attraction of the formulation. However, in the spirit of a system with a maximum, it is worth noting what happens when there are two classes of people, one whose earnings never rise above the Social Security maximum, and one whose earnings never drop below the Social Security maximum. Equation (20) still gives after-tax income for the poor group, and inequality among them spreads as in the previous section. For the rich group, after-tax income is

$$y_{it} = (1 - \tau)(\mu_i + \varepsilon_{it}) + \tau(\mu_i + \varepsilon_{it} - m) + \frac{(\tau\bar{\mu}_1 + \tau m)}{2}, \quad (23)$$

where m is the Social Security maximum, $\bar{\mu}_1$ is mean earnings of the poorer group, and we have assumed that there are equal numbers in the two groups. (The first term is what is left if tax was paid on everything, the second term is the rebate of tax above the maximum, and the last term is the shared benefit.) Equation (23) can be rewritten as

$$y_{it} = \mu_i - \frac{\tau(m - \bar{\mu}_1)}{2} + \varepsilon_{it}, \quad (24)$$

which makes the straightforward point that those above the maximum no longer participate in the risk sharing, only in the redistribution. As a result, the Social Security system with the two groups will limit the rate of spread of inequality among the poorer group, but not among the richer group, although it will bring the two groups closer together than they would have been in the absence of the system.

4.2.5 Finite Lives with Retirement

With finitely lived consumers we can have a more realistic Social Security system, in which the taxes are repaid in retirement rather than instantaneously. One point to note about retirement is that it induces a fall in earnings at the time of retirement, a fall that enters into the determination of saving (see equation [6]). When there is a unit root in earnings, earnings immediately prior to retirement have a unit root, and so does the drop in earnings at retirement. In consequence, saving, which must cover this drop in earnings, is no longer stationary but integrated of order one, so that assets, which are cumulated saving, are integrated of order two. The spread of inequality in assets is therefore an order of integration faster than the spread of inequality in consumption and disposable income. However, this seems more a matter of degree than an essential difference.

People work until age R and die at age T. The consumption innovation formula is only slightly different:

$$\beta_t \Delta c_t = \eta_t = \frac{r}{1 + r} \sum_{k=0}^{R-t} \frac{1}{(1 + r)^k} (E_t - E_{t-1}) y_{t+k}, \quad (25)$$

where the annuity factor β_t is given by

$$\beta_t \equiv 1 - \frac{1}{(1+r)^{(T-t+1)}} = (1+r)\beta_{t-1} - r. \tag{26}$$

From equation (25), we can write

$$c_t = c_0 + \sum_{s=0}^{t} \beta_s^{-1} \eta_t. \tag{27}$$

Hence, in the i.i.d. case previously considered,

$$\operatorname{var}_t(c) = \operatorname{var}_0(c) + \frac{r^2}{(1+r)^2} \sigma_z^2 \sum_{s=0}^{t} \beta_s^{-2}. \tag{28}$$

With the Social Security scheme, after-tax earnings while working is

$$\begin{aligned} y_{it} &= (1-\tau)(\mu_i + \varepsilon_{it}) \\ &= (1-\tau)\mu_i + (1-\tau)(w_t + z_{it}). \end{aligned} \tag{29}$$

With a uniform distribution of ages, the benefits while retired in year $R + s$ are

$$\frac{R\tau(\bar{\mu} + w_{R+s})}{T - R}. \tag{30}$$

With certainty equivalence, only the expected present value of this matters (which is a constant given the i.i.d. assumption) so that, once again, although the levels of consumption are altered, there is no change to the innovation of consumption, nor to the rate at which the various inequalities spread.

These results would clearly be different with either an autocorrelation structure of the macro component of earnings such that current innovations had information about what will happen in retirement, although this issue seems hardly worth worrying about; or precautionary motives or borrowing restrictions, such that transactions that leave net present value unaffected can have real effects on the level and profile of consumption. Without quadratic preferences, and without the ability to borrow, we cannot even guarantee the basic result that uncertainty in earnings causes consumption and income inequality to increase with age. In consequence, we have little choice but to specify a model and to simulate the effects of alternative Social Security policies, and this is the topic of section 4.3. Of course, it might reasonably be argued that the purpose of Social Security is not well captured within any of these models, and that present-value neutral “forced” saving has real effects, not because of precautionary motives or borrowing restrictions, but because people are myopic or otherwise unable to make sensible retirement plans on their own. We are sympathetic to the general argument, but have nothing to say about such a case; without a more explicit model of behavior, it is not possible to conclude anything about the effects of Social Security reform on inequality.

4.3 Social Security with Precautionary Saving or Borrowing Constraints

4.3.1 Describing the Social Security System

When consumers cannot borrow, or when they have precautionary motives for saving, the timing of income affects their behavior. In consequence, we need to be more precise about the specification of the Social Security system and its financing. We assume that there is a constant rate of Social Security tax on earnings during the working life, levied at rate τ, and that during retirement, the system pays a two-part benefit. The first part, G, is a guaranteed floor that is paid to everyone, irrespective of their earnings or contribution record. The second part, V_i, is individual-specific and depends on the present value of earnings (or contributions) over the working life. We write S_i for the annual payment to individual i after retirement, so that

$$S_i = G + V_i = G + \tilde{\alpha}\sum_{j=1}^{R-1} y_{ij}^b(1 + r)^{R-j} \tag{31}$$
$$= G + \alpha\sum_{j=1}^{R-1} y_{ij}(1 + r)^{R-j},$$

where $\alpha = \tilde{\alpha}/(1 - \tau)$. The size of the parameter α determines the extent of the link between earnings in work and Social Security payments in retirement. When we consider the effects of different Social Security systems on inequality, we shall consider variations in α and G while holding the tax rate τ constant. As we shall see below, this is equivalent to devoting a larger or smaller share of Social Security tax revenues to individual accounts. When α is high relative to G (personal saving accounts), the system is relatively autarkic, and there is relatively little sharing of risk. Conversely, when G is large and α small (the current system), risk sharing is more important, and we expect inequality to be lower.

The government finances the Social Security system in such a way as to balance the budget in present-value terms within each cohort. If we use the date of retirement as the base for discounting, the present value of government revenues from the Social Security taxes levied on the cohort about to retire is given by

$$\tau\sum_{j=1}^{R-1}\sum_{i=1}^{N} y_{ij}^b(1 + r)^{R-j} = \tau\sum_{j=1}^{R-1} Y_t(1 + r)^{R-j}, \tag{32}$$

where N is the number of people and Y_t is aggregate before-tax earnings for the cohort in year t. This must equal the present value at R of Social Security payments, which is

$$\sum_{i=1}^{N}\sum_{j=R}^{T}(1 + r)^{R-j}\left[G + \tilde{\alpha}\sum_{j=1}^{R-1} y_{ij}^{b}(1 + r)^{R-j}\right]. \tag{33}$$

The budget constraint that revenues equal outlays, that equation (32) equals equation (33), gives a relationship between the three parameters of the Social Security system, τ, G, and $\tilde{\alpha}$, namely

$$G + \tilde{\alpha}\bar{y}^{*} = \frac{\tau\bar{y}^{*}}{\sum_{j=R}^{T}(1 + r)^{R-j}}, \tag{34}$$

where $\bar{y}^{*}$ is the average over all consumers of the present value of lifetime earnings,

$$\bar{y}^{*} = \left(\frac{1}{N}\right)\sum_{j=1}^{R-1} Y_j(1 + r)^{R-j}. \tag{35}$$

Equation (34) tells us that we can choose any two of the three parameters, G, α (or $\tilde{\alpha}$), and τ, and what is implied for the third. It also makes clear that, after appropriate scaling, and holding the guarantee fixed, increases in α—the earnings-related or autarkic part of the system—are equivalent to increases in the rate of the Social Security tax, given that the government is maintaining within-cohort budget balance.

The link between earnings-related Social Security payments and individual accounts can be seen more clearly if we reparameterize the system. Suppose that V_i, the earnings-related component of the Social Security payment, is funded out of a fraction of Social Security taxes set aside for the purpose, or equivalently, that a fraction φ of the tax is used to build a personal account, the value of which is used to buy an annuity at retirement. Equating the present value of each annuity V_i to the present value of contributions gives the relationship between α and φ,

$$\varphi = \left(\frac{\tilde{\alpha}}{\tau}\right)\sum_{j=R}^{T}(1 + r)^{R-j}. \tag{36}$$

Hence, any increase in the earnings-related component of Social Security through an increase in α (or $\tilde{\alpha}$) can be thought of as an increase in the fraction of Social Security taxes sequestered into personal accounts. Equation (34), which constrains the parameters of the Social Security system, can be rewritten in terms of φ as

$$G = \frac{\bar{y}^{*}\tau(1 - \varphi)}{\sum_{j=R}^{T}(1 + r)^{j-R}}. \tag{37}$$

Note also that the individual Social Security payment in equation (31) can be rewritten as

$$(38)\qquad S_i = \frac{\tau}{\sum_{j=R}^{T}(1+r)^{R-j}}\left[(1-\varphi)\bar{y}^* + \varphi\sum_{j=1}^{R-1} y_{ij}^{b}(1+r)^{R-j}\right],$$

so that each person's Social Security benefits are related to a weighted average of their own lifetime earnings and the average lifetime earnings of their entire cohort.

If the above scheme were implemented for permanent-income consumers who are allowed to borrow and lend at will, the component of Social Security taxes that goes into personal accounts would have no effect on individual consumption nor, therefore, on its distribution across individuals. Although the scheme forces people to save, it is fair in present value terms, and so has no effect on the present value of each individual's lifetime resources. Moreover, although taxes are paid now and benefits received later, such a transfer can be undone by appropriate borrowing and lending. If the Social Security tax rate is τ, and a fraction φ is invested in a personal account, it is as if the tax rate were reduced to $\tau(1 - \varphi)$, and the rate of increase in the consumption and income variance will be higher. Of course, none of these results hold if consumers are not allowed to borrow, or if preferences are other than quadratic.

4.3.2 Modeling Consumption Behavior

Although we shall also present results from the permanent income hypothesis, our preferred model is one with precautionary motives based on that in Gourinchas and Parker (2002) and Ludvigson and Paxson (2001), with the addition of retirement and a simple Social Security system. The specification and parameters are chosen to provide a reasonable approximation to actual behavior so that, even though it is not possible to derive closed-form solutions for the results, we can use simulations to give us some idea of the effects of the reforms.

Consumers have intertemporally additive isoelastic utility functions and, as before, they work through years 1 to $R - 1$, retiring in period R and dying in period T. The real interest rate is fixed, but the rate of time preference δ is (in general) different from r, so that consumers satisfy the familiar Euler equation

$$(39)\qquad c_t^{-\rho} = \beta(1+r)E_t(c_{t+1}^{-\rho}),$$

where ρ is the inverse of the intertemporal elasticity of substitution and $\beta = (1 + \delta)^{-1}$. After-tax earnings, where taxes include Social Security taxes, evolves according to the (also fairly standard) nonstationary process

$$(40)\qquad \ln y_t = \ln y_{t-1} + \gamma + \varepsilon_t - \lambda\varepsilon_{t-1},$$

which derives from a specification in which log earnings are the sum of a random walk with drift γ and white noise transitory earnings. The quantity λ is the parameter of the moving average process for the change in earnings and is an increasing function of the ratio of the variances of the transitory and random walk components, respectively. Consumers are assumed to be unable to borrow, which requires a modification of equation (39) (see below). One reason for this assumption is to mimic the United States, where it is illegal to borrow against prospective Social Security income. A second reason is to rule out the possibility that people borrow very large sums early in life to finance a declining consumption path over the life cycle. This prohibition could be enforced in other ways, such as the "voluntary" borrowing constraints in Carroll (1997) that result from isoelastic utility coupled with a finite probability of zero earnings. We do not find Carroll's income process empirically plausible, and it seems simpler to rule out borrowing explicitly rather than to choose the form of the earnings process to do so. Our calculations for the permanent income case are made with and without borrowing constraints, which will give some idea of the effects of the borrowing constraints in the other models.

Our procedure is as follows. Given values for the real interest rate, the rate of time preference, the intertemporal elasticity of substitution, the moving average parameter in income growth, and two out of three parameters of the Social Security system, we calculate a set of policy functions for each year of a forty-year working life. After retirement, there is no further uncertainty, and consumption can be solved analytically for each of the twenty years remaining. We assume that the Social Security system presented in section 4.3.2 has been in place for a long time, that its parameters are fixed, and that people understand how it works, including the government's intertemporal budget constraint. In particular, they understand the implications of innovations to their earnings for the value of their annuities in retirement. We do not require consumers to take into account the effects of successive macroeconomic shocks on the size of the Social Security guarantee G. Instead, we assume that the government sets G to the value that satisfies the budget constraint in expectation for each cohort, and that deficits and surpluses from cumulated macro shocks are passed on to future generations. There are, however, no macro shocks in the simulations reported below.

In each period of the working life, the ratio of consumption to earnings can be written as function of three state variables. These are defined as follows. Define cash on hand $x_t = A_t + y_t$, which, by equation (1), evolves during the working life $t < R$ according to

$$x_t = (1 + r)(x_{t-1} - c_{t-1}) + y_t \, . \tag{41}$$

During retirement, for $t \geq R$,

$$x_t = (1 + r)(x_{t-1} - c_{t-1}) + S. \tag{42}$$

If w_t is the ratio of cash on hand to earnings, and θ_t the ratio of consumption to earnings, then equation (34) becomes, for $t < R$,

$$w_t = \frac{(1 + r)(w_{t-1} - \theta_{t-1})}{g_t} + 1, \tag{43}$$

where g_t is the ratio of current to lagged income, y_t/y_{t-1}. To derive corresponding equations for the dynamics of Social Security, define S_t as the current present value of the annual Social Security payment to which the consumer would be entitled if he or she earned no more income between year t and retirement. Hence, for $t < R$,

$$S_t = G(1 + r)^{-(R-t)} + \alpha\sum_{j=1}^{t} y_j(1 + r)^{t-j}, \tag{44}$$

while for $t \geq R$, S_t is constant and given by equation (31). Noting that earnings in year R is zero, equation (44) satisfies, for $t \leq R$,

$$S_t = (1 + r)S_{t-1} + \alpha Y_t \tag{45}$$

and is constant thereafter. If we define σ_t, the ratio of S_t to current earnings and thus the "Social Security replacement rate," the corresponding evolution equation is

$$\sigma_t = \frac{S_t}{y_t} = (1 + r)\frac{\sigma_{t-1}}{g_t} + \alpha. \tag{46}$$

With borrowing constraints, which imply that consumption cannot be greater than cash on hand, or that the consumption ratio be no larger than the cash on hand ratio, the Euler equation (39) is modified to

$$\theta_t^{-\rho} = \max\left[\beta(1 + r)E_t(g_{t+1}^{-\rho}\theta_{t+1}^{-\rho}), w_t^{-\rho}\right]. \tag{47}$$

We write the consumption ratio θ_t as a function of the cash on hand ratio w_t, the Social Security replacement rate σ_t, and the current innovation to earnings ε_t (which is required because, with positive λ, high earnings growth in one period predicts low earnings growth in the next), and then use equation (47) to solve backward for the policy function in each period, starting from the closed-form solution for consumption in the first year of retirement.

Armed with the policy functions, we simulate lifetime stochastic earnings profiles for each of 1,000 people. The logarithm of initial earnings is drawn from a normal distribution with mean ln (20,000) and a standard deviation of 0.65, the latter chosen to give an initial Gini coefficient that roughly corresponds to what we see in the data from the CPS. The drift (expected rate of growth) of earnings is set at 2 percent a year. For any given value of the replacement parameter α and the Social Security tax rate τ, the corresponding value of the Social Security guarantee G is set from equation (34) using actual realized earnings, which, as we have already noted, is potentially problematic if macro shocks are important. The value of G also gives the initial value of σ_t at the beginning of life. The calculated policy functions are then used to simulate life-cycle consumption for each of the 1,000 people, and these trajectories are used to assess lifetime inequality as a function of the design of the Social Security system. Different simulations use the same 1,000 sets of earnings realizations, so that comparisons across Social Security regimes reflect the regime parameters and not the specific draws.

4.3.3 Social Security Design and Inequality: Results with Constant Interest Rates

The model is solved under the following assumptions. The interest rate r is set at 3 percent, and the rate of time preference at either 3 or 5 percent. The drift of the earnings process is set at 2 percent a year, the moving average parameter λ to 0.4, and the standard deviation of the innovation (in logs) to be 0.25. The coefficient of relative risk aversion is set to 3, so that the intertemporal elasticity of substitution is one-third. We also include a certainty equivalent case, with and without borrowing restrictions, in which the rate of interest is set equal to the rate of time preference at 3 percent. There are four cases carried through the analysis: (1) isoelastic preferences, no borrowing, $r = 0.03$, $\delta = 0.05$; (2) isoelastic preferences, no borrowing, $r = 0.03$, $\delta = 0.03$; (3) quadratic preferences, no borrowing, $r = 0.03$, $\delta = 0.03$; and (4), quadratic preferences, borrowing allowed, $r = 0.03$, $\delta = 0.03$. The Social Security tax rate is set at its current value of 12.4 percent of before-tax earnings and there are no other taxes or benefits. The Social Security systems we consider are indexed on the level of the Social Security guarantee G, which takes the values ($0, $5,000, $10,000, $15,000, $20,000); given the tax rate, these values translate into corresponding values for α or, perhaps more revealingly, into values for φ, the share of the tax devoted to personal accounts (1, 0.811, 0.623, 0.434, 0.245). These different sets of parameters have quite different implications for the dispersion in Social Security payments among retirees. For example, our simulation results indicate that with a guarantee of $0, the person at the 10th percentile (ranked by the present value of lifetime earnings) receives an annual Social Security payment of $6,405, in contrast to a payment of $52,639 for the person at the 90th percentile. When the guarantee is increased to $20,000, this spread declines to $21,569 for the 10th percentile, and to $32,896 for the 90th.

Figure 4.1 shows the averages over the 1,000 consumers of the simulated

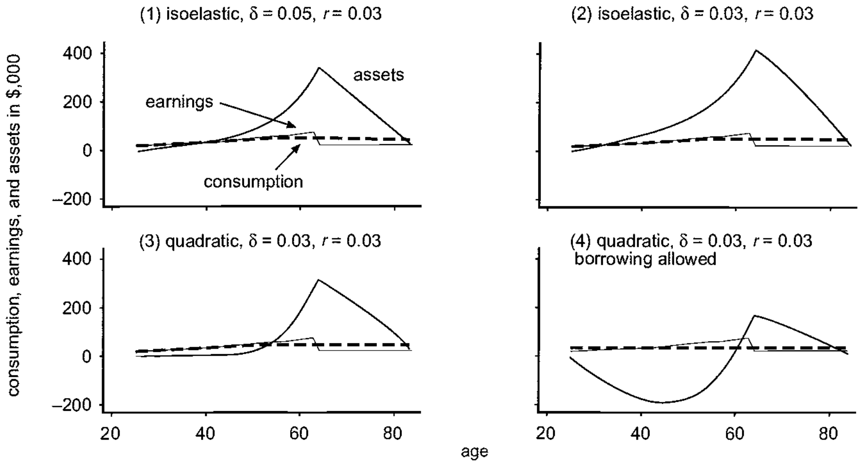

Fig. 4.1 Age profiles of consumption, earnings (inclusive of transfers), and assets for different specifications, G = \$5,000

trajectories of income (earnings prior to retirement and receipts from Social Security after retirement), consumption, and cash on hand (earnings plus assets excluding Social Security assets) for the four models all with G set at \$5,000. These graphs are shown to demonstrate that the various models do indeed generate standard life-cycle profiles. Earnings is the same in each of the three graphs. Consumption is flat over the life cycle in the certainty equivalent case when borrowing is allowed, but rises in the models with precautionary motives and borrowing constraints, and in the quadratic case with borrowing constraints. Indeed, the quadratic case with no borrowing (on the bottom left) and the isoelastic "impatient" case with no borrowing (top left) generate similar profiles. With more patient (lower δ) consumers in the top right panel, there is more accumulation during the working life, and assets prior to retirement are higher. The certainty equivalent consumers in the bottom right panel have expectations of earnings growth and so engage in substantial borrowing early in life but, even so, have some net assets prior to retirement.

Figure 4.2 shows the average consumption profiles for the four different models (in the four panels, as before) and for the five different Social Security schemes (in each panel). To a first approximation, and with the tax rate held fixed, the choice of system has no effect on the lifetime profile of consumption. Figure 4.2 also shows more clearly than figure 4.1 the lifetime shape of consumption in the four models: Precautionary motives or borrowing restrictions drive the increase in consumption over the working period; in the top left panel, where impatience is greater than the interest rate, consumption declines after retirement once all uncertainty is resolved. For the cases with precautionary motives or borrowing restrictions, average consumption during retirement is somewhat higher in the regimes with the higher minimum guarantee. This appears to be a consequence of the borrowing constraints. Those consumers who draw poor earnings throughout their lives, and who would like to borrow against their Social Security but cannot, have higher consumption in retirement when the guarantee becomes available. In effect, such consumers are being forced to save for higher consumption in retirement than they would choose if left to themselves. Such effects are absent in the pure certainty equivalent case where borrowing is allowed.

Figure 4.3 plots the Gini coefficients of consumption by age and shows how consumption inequality evolves in the various models and for the different Social Security systems. The Gini coefficients, together with interquartile ranges of the logarithm of consumption, are given in numerical form in table 4.1. In all of the models, consumption inequality is higher at all ages the lower the Social Security guarantee (the higher the fraction of taxes invested in personal accounts) and the more autarkic the system. A higher guarantee with its associated lower limit to lifetime earnings causes consumption inequality to be lower from the start of the life cycle, though the early effects are strongest in the pure certainty equivalent case, and manifest themselves only later in life in the models with borrowing constraints. With a higher guarantee, and less in individual accounts, the system has more sharing, so that individual earnings innovations have less effect on consumption because the good (or ill) fortune will be shared with others. Although this sharing is implemented only after retirement, because consumption is smoothed over the life cycle, the effect on inequality works at all ages to an extent determined by the assumptions about prefer-

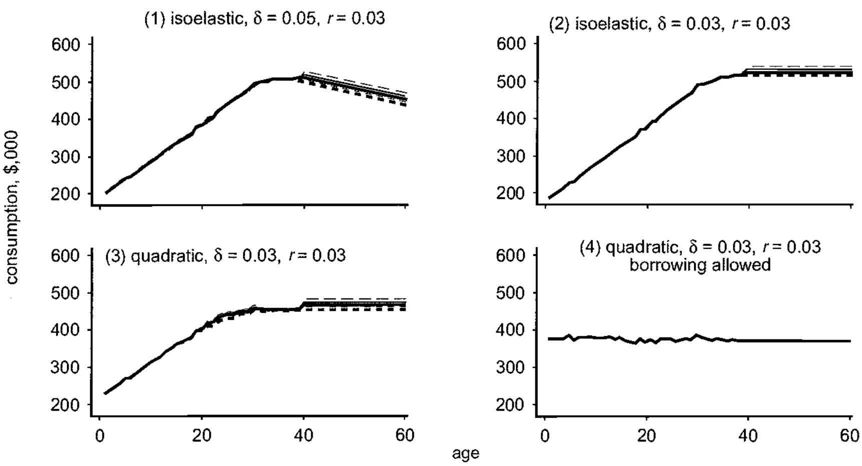

Fig. 4.2 Consumption profiles under different specifications and alternative Social Security rules

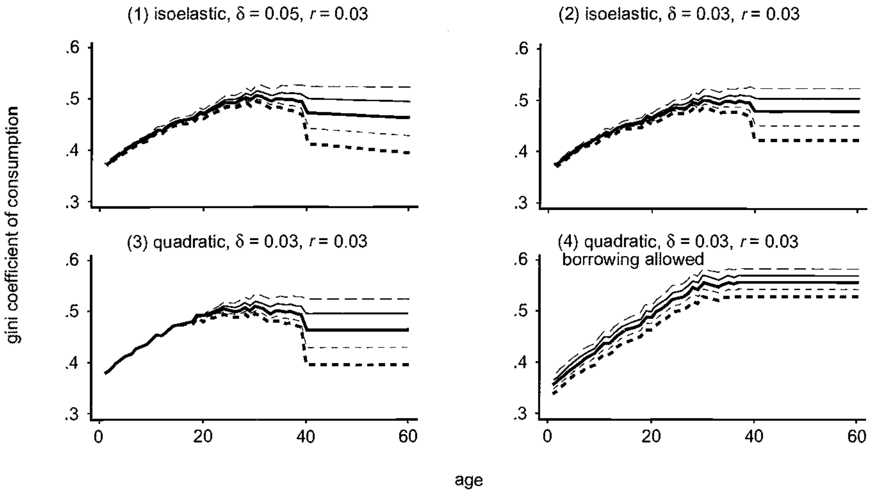

Fig. 4.3 Consumption inequality for different specifications and Social Security systems

ences, growth, and borrowing constraints. When borrowing constraints are imposed in an environment with earnings growth, consumption smoothing is inhibited, and the effects of risk sharing on inequality are more apparent in the later than in the earlier phases of the life cycle. These results are not sensitive to the choice of inequality measure. The interquartile ranges, although somewhat jumpier, display patterns that are similar to the Gini coefficients.

Figure 4.3 also shows a sharp drop in consumption inequality after retirement, particularly when the guarantee in the Social Security system is relatively large. Once again, this comes from the borrowing constraints and the inability of lifetime unlucky consumers to borrow against the Social Security system. These people have very low consumption immediately prior to retirement, which exaggerates inequality. The effect vanishes as Social Security becomes available and their consumption rises. In the cases where the guarantee is large, there is also some decline in inequality prior to retirement. Although no theoretical reason prohibits this, we have not so far developed a convincing explanation of why it should occur. Panel 1 of figure 4.3 also shows a small decline in consumption inequality during retirement. This is due to the combination of borrowing constraints and impatience ($r < \delta$). Unconstrained consumers choose declining consumption paths during retirement (at a constant and common rate of about 0.64 percent per year), while those who are constrained simply consume their constant Social Security income. The result is a compression of the distribution of consumption.

Table 4.1 **Gini Coefficients for Consumption and Interquartile Ranges for Logarithm of Consumption, with Different Social Security Plans**

	G = \$0		G = \$5,000		G = \$10,000		G = \$15,000		G = \$20,000	
Age	Gini	IQR	Gini	IQR	Gini	IQR	Gini	IQR	Gini	IQR
Isoelastic preferences, r = 0.03, δ = 0.05, borrowing constraints										
25	0.37	0.910	0.37	0.905	0.37	0.903	0.37	0.900	0.37	0.894
29	0.41	1.034	0.40	1.024	0.40	1.020	0.40	1.018	0.40	1.011
34	0.44	1.024	0.43	1.011	0.43	1.011	0.43	1.009	0.43	0.998
39	0.46	1.098	0.46	1.079	0.46	1.082	0.45	1.079	0.45	1.064
44	0.48	1.190	0.47	1.177	0.47	1.175	0.47	1.173	0.46	1.154
49	0.50	1.283	0.50	1.249	0.49	1.243	0.49	1.236	0.48	1.230
54	0.52	1.304	0.51	1.240	0.51	1.194	0.50	1.180	0.49	1.176
59	0.52	1.287	0.51	1.199	0.50	1.122	0.49	1.056	0.48	1.035
64	0.52	1.292	0.50	1.204	0.47	1.098	0.44	0.993	0.41	0.896
69	0.52	1.292	0.50	1.204	0.47	1.093	0.44	0.986	0.41	0.876
74	0.52	1.292	0.50	1.201	0.47	1.089	0.44	0.976	0.40	0.859
79	0.52	1.292	0.50	1.191	0.46	1.086	0.43	0.956	0.40	0.835
84	0.52	1.288	0.49	1.189	0.46	1.066	0.43	0.934	0.39	0.812
Isoelastic preferences, r = 0.03, δ = 0.03, borrowing constraints										
25	0.37	0.899	0.37	0.888	0.37	0.880	0.36	0.876	0.36	0.872
29	0.40	1.033	0.40	1.016	0.40	1.010	0.40	1.007	0.39	1.000
34	0.43	1.011	0.43	0.992	0.43	0.984	0.42	0.982	0.42	0.970
39	0.46	1.092	0.45	1.068	0.45	1.054	0.44	1.049	0.44	1.039
44	0.47	1.164	0.46	1.136	0.46	1.123	0.46	1.118	0.45	1.104
49	0.50	1.238	0.49	1.205	0.48	1.181	0.48	1.173	0.47	1.168
54	0.52	1.283	0.51	1.216	0.50	1.161	0.49	1.127	0.48	1.113
59	0.52	1.275	0.51	1.191	0.49	1.115	0.48	1.039	0.48	0.985
64	0.52	1.284	0.50	1.207	0.47	1.130	0.45	1.033	0.42	0.931
69	0.52	1.284	0.50	1.207	0.47	1.130	0.45	1.033	0.42	0.931
74	0.52	1.284	0.50	1.207	0.47	1.130	0.45	1.033	0.42	0.931
79	0.52	1.284	0.50	1.207	0.47	1.130	0.45	1.033	0.42	0.931
84	0.52	1.284	0.50	1.207	0.47	1.130	0.45	1.033	0.42	0.931

				Quadratic preferences, r = 0.03, δ = 0.03, borrowing constraints						
25	0.38	0.918	0.38	0.918	0.38	0.918	0.38	0.918	0.38	0.918
29	0.41	1.050	0.41	1.050	0.41	1.050	0.41	1.050	0.41	1.050
34	0.44	1.067	0.44	1.067	0.44	1.067	0.44	1.067	0.44	1.067
39	0.47	1.167	0.47	1.167	0.47	1.167	0.47	1.166	0.47	1.165
44	0.49	1.231	0.49	1.229	0.48	1.227	0.48	1.224	0.48	1.218
49	0.51	1.327	0.51	1.296	0.50	1.291	0.49	1.283	0.49	1.276
54	0.53	1.315	0.52	1.233	0.51	1.177	0.50	1.157	0.49	1.152
59	0.53	1.287	0.51	1.184	0.50	1.093	0.49	1.045	0.48	1.009
64	0.52	1.298	0.49	1.183	0.46	1.050	0.43	0.929	0.39	0.815
69	0.52	1.298	0.49	1.183	0.46	1.050	0.43	0.929	0.39	0.815
74	0.52	1.298	0.49	1.183	0.46	1.050	0.43	0.929	0.39	0.815
79	0.52	1.298	0.49	1.183	0.46	1.050	0.43	0.929	0.39	0.815
84	0.52	1.298	0.49	1.183	0.46	1.050	0.43	0.929	0.39	0.815
				PIH: Quadratic preferences, r = 0.03, δ = 0.03, no borrowing constraints						
25	0.37	0.899	0.36	0.872	0.36	0.845	0.35	0.818	0.34	0.792
29	0.41	1.029	0.40	0.996	0.39	0.965	0.38	0.934	0.37	0.903
34	0.45	1.061	0.44	1.025	0.43	0.990	0.42	0.955	0.41	0.922
39	0.48	1.193	0.47	1.149	0.46	1.107	0.45	1.066	0.44	1.026
44	0.51	1.333	0.50	1.280	0.48	1.229	0.47	1.180	0.46	1.133
49	0.55	1.455	0.54	1.389	0.52	1.328	0.51	1.270	0.50	1.215
54	0.58	1.515	0.57	1.442	0.55	1.377	0.54	1.316	0.53	1.258
59	0.58	1.526	0.57	1.457	0.55	1.388	0.54	1.323	0.53	1.262
64	0.58	1.560	0.57	1.493	0.55	1.422	0.54	1.356	0.53	1.294
69	0.58	1.560	0.57	1.493	0.55	1.422	0.54	1.356	0.53	1.294
74	0.58	1.560	0.57	1.493	0.55	1.422	0.54	1.356	0.53	1.294
79	0.58	1.560	0.57	1.493	0.55	1.422	0.54	1.356	0.53	1.294
84	0.58	1.560	0.57	1.493	0.55	1.422	0.54	1.356	0.53	1.294

Source: Authors' calculations.

Notes: "Gini" refers to the Gini coefficient for consumption. IQR is the interquartile range of the logarithm of consumption. PIH is the permanent income hypothesis.

Overall, the results in figure 4.3 and table 4.1 show that as we move from one extreme to the other, from putting everything into individual accounts and giving no guarantee (a Social Security system than confines itself to compulsory saving) to a guaranteed floor of $20,000 with only one-fourth of Social Security taxes going to personal accounts, the Gini coefficient of consumption increases by between 5 and 6 percentage points on average over the life cycle, less among the young, and more among the old. This is a large increase, exceeding the increase in consumption inequality in the United States during the inequality boom from the early to the mid-1980s. For example, the Gini coefficient of total consumption for urban households from the U.S. Consumer Expenditure Survey rose from 0.37 in 1981 to 0.41 in 1986.

Table 4.2 shows "poverty rates" by age for the different models and Social Security systems. An individual is defined as being in poverty if annual consumption is less than $10,000. This poverty threshold was arbitrarily chosen, but it delivers total poverty rates that are not very different from those in the United States. For example, with G equal to $5,000, the total poverty rate is 12.5 percent for the first model. We are more concerned with how poverty varies with age than with its level. The age profiles of poverty are similar for the first three models, in which there are borrowing constraints. Poverty rates decline up to retirement age: Constrained consumers are more likely to be poor when they are young, and earnings are low. Poverty in retirement depends on the value of the Social Security guarantee. When the guarantee is greater than or equal to the poverty threshold, poverty in retirement must equal zero. For smaller values of the guarantee, the poverty rate in retirement is generally less than during working years. However, in one case—that of isoelastic preferences and $r < \delta$—the poverty rate grows during retirement. In this case, impatient consumers reduce consumption over time, and increasingly fall below the threshold.

The fourth model, with quadratic preferences and no borrowing constraints, yields very different results. Poverty rates increase with age up to retirement. Average consumption is constant over the life cycle, and the increasing dispersion in consumption with age implies that consumers will increasingly fall below the threshold. Increases in the poverty rate cease at retirement. However, Social Security guarantees in excess of the poverty threshold do not eliminate poverty, since (in this model) individuals are free to borrow against the guarantee during working years. Higher Social Security guarantees do, in fact, reduce poverty, but they do so at all ages, by making lifetime wealth more equal across individuals.

Figure 4.4 compares our simulated patterns of inequality over the life cycle with those calculated from the data in the CEX and reported in Deaton and Paxson (1994). By construction, the life-cycle profile of simulated earnings inequality is similar to the actual profile. Simulated consumption

Table 4.2 **Poverty Rates (fraction of age group with consumption less than \$10,000) with Different Social Security Plans**

	G				
Age	\$0	\$5,000	\$10,000	\$15,000	\$20,000
		Isoelastic preferences, $r = 0.03$, $\delta = 0.05$			
25	0.239	0.238	0.238	0.237	0.238
29	0.210	0.209	0.207	0.207	0.207
34	0.177	0.176	0.175	0.175	0.175
39	0.158	0.153	0.153	0.153	0.153
44	0.146	0.140	0.143	0.142	0.139
49	0.133	0.127	0.127	0.128	0.128
54	0.120	0.115	0.116	0.117	0.117
59	0.119	0.104	0.103	0.106	0.105
64	0.119	0.061	0.000	0.000	0.000
69	0.123	0.069	0.000	0.000	0.000
74	0.134	0.073	0.000	0.000	0.000
79	0.141	0.076	0.000	0.000	0.000
84	0.144	0.081	0.000	0.000	0.000
		Isoelastic preferences, $r = 0.03$, $\delta = 0.03$			
25	0.264	0.262	0.258	0.258	0.258
29	0.228	0.227	0.225	0.225	0.225
34	0.186	0.185	0.179	0.178	0.178
39	0.155	0.154	0.153	0.150	0.150
44	0.141	0.135	0.134	0.134	0.134
49	0.129	0.115	0.115	0.115	0.114
54	0.118	0.106	0.099	0.098	0.098
59	0.112	0.095	0.091	0.093	0.094
64	0.113	0.067	0.000	0.000	0.000
69	0.113	0.067	0.000	0.000	0.000
74	0.113	0.067	0.000	0.000	0.000
79	0.113	0.067	0.000	0.000	0.000
84	0.113	0.067	0.000	0.000	0.000
		Quadratic preferences, $r = 0.03$, $\delta = 0.03$			
25	0.195	0.195	0.195	0.195	0.195
29	0.186	0.186	0.186	0.186	0.186
34	0.163	0.163	0.163	0.163	0.163
39	0.152	0.152	0.152	0.152	0.152
44	0.149	0.149	0.149	0.149	0.149
49	0.144	0.141	0.138	0.139	0.139
54	0.138	0.129	0.125	0.125	0.126
59	0.147	0.126	0.118	0.123	0.124
64	0.135	0.073	0.000	0.000	0.000
69	0.135	0.073	0.000	0.000	0.000
74	0.135	0.073	0.000	0.000	0.000
79	0.135	0.073	0.000	0.000	0.000
84	0.135	0.073	0.000	0.000	0.000

(*continued*)

Table 4.2 (continued)

	G				
Age	\$0	\$5,000	\$10,000	\$15,000	\$20,000
	Quadratic preferences, r = 0.03, δ = 0.03, no borrowing constraints				
25	0.055	0.047	0.036	0.030	0.014
29	0.082	0.068	0.054	0.043	0.030
34	0.116	0.100	0.085	0.068	0.051
39	0.153	0.143	0.125	0.114	0.097
44	0.202	0.187	0.175	0.149	0.134
49	0.247	0.224	0.203	0.183	0.163
54	0.260	0.242	0.221	0.194	0.170
59	0.277	0.259	0.234	0.209	0.192
64	0.280	0.264	0.239	0.212	0.197
69	0.280	0.264	0.239	0.212	0.197
74	0.280	0.264	0.239	0.212	0.197
79	0.280	0.264	0.239	0.212	0.197
84	0.280	0.264	0.239	0.212	0.197

Source: Authors' calculations.

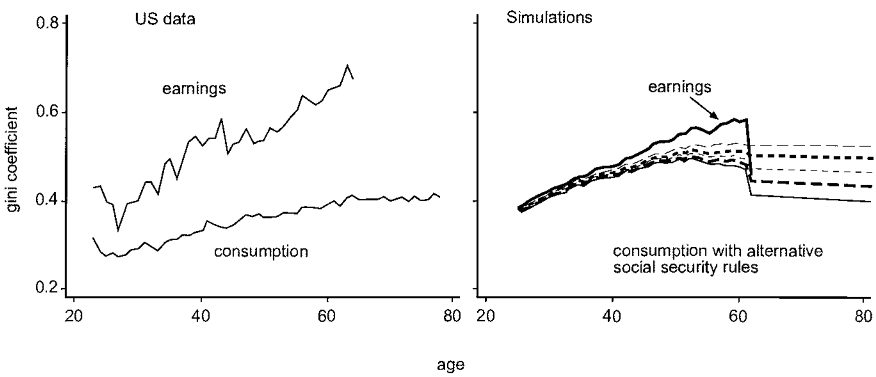

Fig. 4.4 Actual and simulated inequality of earnings and consumption

inequality (from the "impatient" isoelastic case) is too high relative to the actuals; perhaps the borrowing restrictions are preventing consumption from being sufficiently smoothed. Nevertheless, the upward drift of consumption inequality is very much the same in the data as in the simulations, which also show the effects on inequality of the different Social Security designs.

In figures 4.5 and 4.6 we turn to the life-cycle pattern of inequality in assets, in figure 4.5 for assets excluding Social Security wealth, and in figure 4.6 including Social Security wealth. The permanent income model is excluded from these comparisons because average wealth is negative for much of the life cycle. Total wealth at any given age is defined as the sum

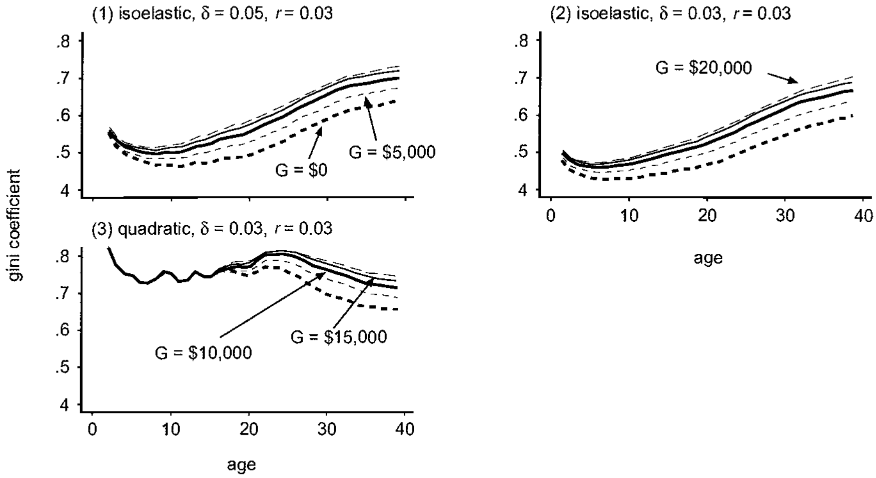

Fig. 4.5 Gini coefficients for assets, excluding Social Security assets

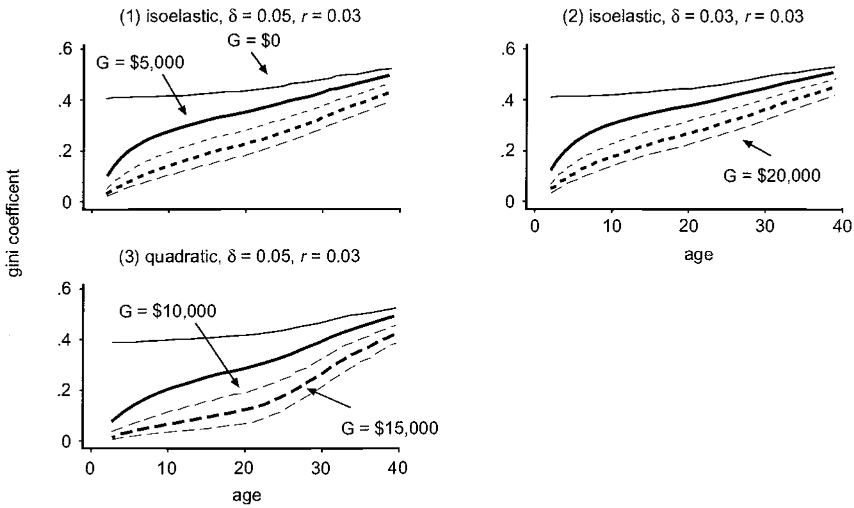

Fig. 4.6 Gini coefficients for total assets, including Social Security assets

of non–Social Security assets A_t and the present discounted value at t of receiving G from retirement R to death T, plus the accumulated balance in the personal saving account, if any. Making the Social Security system more autarkic by holding the Social Security tax constant and devoting more of the revenue to personal accounts and less to a universal guarantee has the opposite effect on inequality of non–Social Security wealth than it does on the inequality of consumption. This is because of the substitutability between saving for retirement inside and outside the Social Security system. If we examine the profiles of asset accumulation by age (not shown here), average non–Social Security accumulations are larger the smaller is the fraction of the Social Security tax invested in individual accounts. There is a similar substitutability in asset inequality; when there is a large Social Security floor for everyone, the resulting equality is partially offset by inequality in private accumulation.

The different patterns of asset inequality for the quadratic case in the bottom left panel, as opposed to the isoelastic cases in the top panels, are associated with the fact that a substantial fraction of the quadratic consumers are credit constrained up until around age forty, so that inequality is high at early ages because so many consumers have exactly nothing. The offsetting of private wealth against Social Security wealth only shows up once the majority of consumers are accumulating private assets, at which point they are no longer credit constrained. In the cases with isoelastic preferences, the borrowing constraints are binding for only a small fraction of young consumers; the variability of earnings and the convexity of marginal utility is enough to overcome impatience and the expected growth of earnings.

When we come to figure 4.6, which shows the inequality of all assets, we see the "standard" pattern restored; the more autarkic the system and the larger the fraction of Social Security taxes devoted to private accounts, the larger is the inequality of assets. Note that the Gini coefficients for *all* assets are much lower than those for private assets; even with personal accounts, the addition of Social Security to private wealth makes the distribution of wealth much more equal. As with consumption, asset inequality rises with age, but does so most rapidly in the cases where insurance is greatest, so that the differences in asset inequalities across the various schemes diminish with age. Even so, the most autarkic systems are the most unequal at all ages.

4.3.4 Social Security Design and Inequality: Results with Variable Interest Rates

The results on asset inequality, and to a lesser extent those on consumption inequality, are likely to be seriously affected by our assumption that everyone earns the same rate of return on their assets. Under some of the early proposals for reform—for example, those from the largest group in the Gramlich report—one of the great virtues of personal accounts was seen as the freedom given to individual consumers to choose their own portfolios. More recent proposals have tended to favor severe restrictions on portfolio choice, perhaps restricting consumers to a limited menu of approved funds, which themselves must adhere to strict portfolio rules. Clearly, allowing different people to obtain different returns adds a new source of inequality, in both assets and in consumption. If, for example, the funds for the minimum guarantee *G* were invested in a common fund at rate *r*, as above, but the personal accounts obtained different rates of return for different individuals, either because of their individual portfolio choices or because of differential management fees, then a move to personal accounts can be expected to increase inequality by more than in the calculations presented thus far. Alternatively, if the limited menu of funds offered different risk-return tradeoffs, and if high earners chose higher returns because they are less risk averse, the availability of the menu would likely translate into higher consumption inequality.

It is not obvious how to construct a model with differential asset returns that is both realistic and computationally tractable. We have so far considered only one simple case. Personal accounts are invested in one of eleven mutual funds, and consumers must choose among them at the outset of the working life. The eleven mutual funds have rates of return from 2.5 to 3.5 percent per year. One can think of the funds as having identical (Standard & Poor's 500) portfolios, but management fees range from zero to 1 percentage point; the equilibrium is maintained by differential advertising and reporting services. We allocate our 1,000 consumers randomly to the eleven mutual funds, with equal probability of receiving any one interest rate; this is a conservative procedure, and inequality would presumably be higher if those with higher earnings were more financially sophisticated and systematically chose the no-load funds. We assume that consumers are forced to convert their retirement accounts into annuities at retirement (using the interest rate to which they have been assigned), and also that the Social Security system gives each consumer a guaranteed amount of $5,000 per year after retirement in addition to the annuity.

The results indicate that there is virtually no increase in consumption inequality before retirement, and very little after retirement, associated with assigning different consumers to different fixed interest rates. The top panel of figure 4.7 shows the Gini coefficient for consumption for the cases described above, with dispersion in interest rates, and the case in which all consumers receive the same interest rate of 3.0 percent. This result may not be not surprising, considering that most saving (whether private or through the Social Security system) is done late in life, when income is high, so that those that receive lower interest rates do not have wealth at retirement that is much lower than those with higher interest rates. Consider, for example, a group of 1,000 consumers whose incomes follow the process described above, each of whom pays 12.4 percent in taxes, 81 percent of which is allocated to private Social Security accounts (thereby generating enough government revenue to fund a $5,000 guaranteed payment to each during retirement). If each member of the group receives an interest rate of 2.5 percent, the average private Social Security account balance upon retirement will equal $278,597. This number will be 22.4 percent higher, or $341,014, with an interest rate of 3.5 percent. The percentage difference in total retirement wealth, including the equalizing guarantee of $5,000 per year, is even smaller, and the difference in average consumption

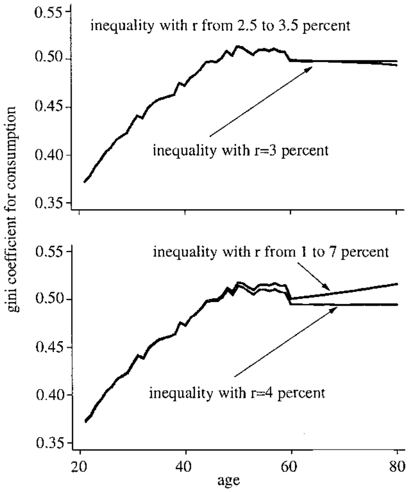

Fig. 4.7 The effects on consumption inequality of a distribution of interest rates

in the first year of retirement for the two interest rates is less than $5,000. This difference is for a spread of one full percentage point; in the exercise conducted above, most consumers have interest rates between the extremes, so there is even less of an effect on overall inequality.

Even with a much wider spread of returns, there is only a modest effect on inequality. The bottom panel of figure 4.7 shows the case in which consumers are distributed over (fixed) rates of return from 1 percent to 7 percent, compared with the case in which all get 4 percent. This can be thought of as the case in which consumers make a choice between equities and bonds at the beginning of their working careers and may never change thereafter. Because the spread is wider, there is more inequality than before, but the effects are modest compared with the other issues examined in this chapter.

It is important to note that assigning consumers to different but fixed rates of interest will not necessarily have the same affects as allowing the interest rate to vary randomly over time for individual consumers. In future work, we plan to examine how interest rate risk, as opposed to interest rate dispersion, affects inequality.

References

Campbell, John Y. 1987. Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis. *Econometrica* 55:1249–73.

Carroll, Christopher. 1997. Buffer-stock saving and the life-cycle permanent income hypothesis. *Quarterly Journal of Economics* 112:1–55.
Deaton, Angus. 1992. *Understanding consumption.* Oxford, England: Clarendon Press.
Deaton, Angus, and Christina Paxson. 1994. Intertemporal choice and inequality. *Journal of Political Economy* 102:437–67.
Feldstein, Martin S. 1998. Is income inequality really a problem? In *Income inequality issues and policy options: A symposium sponsored by the Federal Reserve Bank of Kansas City,* 357–67. N.p.: Federal Reserve Bank of Kansas City.
Flavin, Marjorie. 1981. The adjustment of consumption to changing expectations about future income. *Journal of Political Economy* 89:974–1009.
Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. Consumption over the life-cycle. *Econometrica* (forthcoming).
Ludvigson, Sydney, and Christina H. Paxson. 2001. Approximation bias in linearized Euler equations. *Review of Economics and Statistics* 83:242–56.
Wilkinson, Richard. 1996. *Unhealthy societies: The afflictions of inequality.* London: Routledge.

Comment James M. Poterba

This is an innovative and important chapter that provides new evidence on how Social Security programs affect the distribution of lifetime resources. The chapter presents an elegant analytical treatment of the issues surrounding lifetime inequality and retirement transfer programs. It considers how a somewhat stylized version of the current U.S. Social Security system would affect the degree of inequality in lifetime consumption, and it also explores how a shift toward an individual accounts Social Security system might affect inequality. The chapter provides a very useful starting point for analyzing more complex Social Security arrangements, and many of my comments will focus on potential directions for such extensions. One very attractive feature of the analysis is the presentation of both Gini coefficients for consumption inequality as well as summary statistics for the fraction of the population at different ages that has consumption below a “poverty line” level.

The chapter begins with an elegant treatment of how a simple Social Security system would affect the inequality of consumption, saving, and assets in a stylized economy. The analysis begins with a very general insight. A Social Security system that taxes each worker’s earnings at a fixed

James M. Poterba is the Mitsui Professor of Economics at the Massachusetts Institute of Technology, and both a research associate of and director of the public economics research program at the National Bureau of Economic Research.

rate, and then pays each worker a benefit that is tied to the average earnings level in the population, reduces the variance of consumption and of net-of-tax incomes. If infinitely lived consumers populate the economy, and these consumers save in accordance with the life-cycle hypothesis, then it is possible to derive analytical results for the steady-state variance of consumption with and without the simplified Social Security system.

While the insights from such a model are quite general, the numerical results may fail to describe the impact of actual Social Security programs for two reasons. First, actual Social Security systems do not tax all workers at the same rate, and transfer back a fixed share of economy-wide income. Taxable earnings are often subject to limits, and benefit formulae often incorporate progressive elements that transfer larger amounts, per dollar of taxes paid, to low- than to high-earning individuals. Such programmatic details would be straightforward to incorporate in a somewhat more detailed model of lifetime income and consumption inequality.

The second difficulty with the stylized model in the first part of the chapter is that a substantial body of empirical evidence suggests that many households do not behave in accordance with the simple life-cycle hypothesis. This concern motivates the second part of the chapter, which uses a richer model of consumer behavior, incorporating precautionary motives for saving, to estimate how current and modified Social Security programs could affect consumption inequality. The chapter is careful to consider several different specifications of preferences and to illustrate the sensitivity of key findings to various assumptions. The results are presented both in terms of standard inequality measures, such as the Gini coefficient, and by calculating the fraction of households who experience consumption levels below a prespecified threshold.

The chapter yields several findings of broad interest and importance. First, Social Security systems that levy taxes on realized earnings, but provide benefits that depend in part on aggregate earnings, typically reduce the inequality of lifetime consumption. The current defined benefit (DB) system in the United States has some elements of such a system. Programs of this type reduce inequality in consumption both before and after retirement, and they can have a particularly large impact on the inequality of postretirement consumption when there is a large "guarantee level" that provides a consumption floor for Social Security recipients. With such a guarantee, many households will choose not to save at all for retirement, instead relying on the guarantee level to provide their retirement consumption.

Second, shifting from a DB Social Security system with a guaranteed benefit floor to a system of individual accounts will increase consumption inequality both before and after retirement. The postretirement increase in inequality arises largely from the greater link between preretirement earnings and postretirement resources in individual accounts rather than redistributive defined-benefit systems. The preretirement increase in consumption inequality results in part from the saving adjustments that individuals make in response to the shifting Social Security system. Even if the share of earnings devoted to individual accounts was the same as the share collected by the taxes that finance a DB pension plan, there could be endogenous changes in the saving behavior of individuals and in the resulting pattern of preretirement consumption. This is an important insight, and one that must be considered in future studies of Social Security and inequality.

Third, the numerical results suggest that allowing for differences in the rates of return earned by different individuals within an individual accounts system has a relatively modest impact on the inequality of lifetime consumption. This result seems surprising at first, since one would imagine that greater return variability would result in greater variability of postretirement consumption. While variable rates of return work in this direction, they have a modest effect because most retirement saving is done in the few years before retirement, so the period of time over which differences in returns compound is relatively short. The modest incremental increase in consumption inequality is also, to some extent, a reflection of the very substantial degree of inequality in lifetime income, which translates into heterogeneity in postretirement consumption.

The chapter represents an important start on the very substantial task of modeling how public policies such as Social Security may affect consumption inequality. There are many productive directions in which the current analysis could be extended, however, to provide more information on what might actually happen in a personal accounts retirement system. The remainder of this comment outlines several such directions.

One natural extension is to allow for possible correlation between the rate of return that investors earn on their individual accounts and the level of their lifetime income. There is some evidence from 401(k)-type plans, reported in Poterba and Wise (1998), that higher-income households tend to hold a higher fraction of their 401(k) balances in stocks rather than bonds. Since stocks have historically provided investors with a higher return than bonds, this raises the prospect that those with higher incomes may earn better returns on their individual accounts, on average. Such a correlation would magnify the degree of inequality in retirement resources. It might be particularly important if individuals choose focal values in their asset allocation, such as 0, 50, or 100 percent stock.

A second potentially useful extension would recognize the potential feedback from the Social Security system to the structure of pretax earnings. The current analysis treats the pretax income distribution as given. Yet Social Security systems that tax individuals on their earnings, and return to them a fraction of the economy-wide average earnings at some future date, have incentive effects similar to those of more standard individual income taxes. If redistributive DB Social Security systems are recognized as having different incentive effects than individual accounts systems in which individuals can invest a fraction of their earnings, then it is possible that labor supply would be different under the two systems. If higher tax rates generally discourage labor supply, then one would expect a more compressed distribution of pretax incomes in the DB case than in the individual accounts case. More generally, the possibility that different Social Security systems have different levels of deadweight loss raises a host of additional issues for analysis, because policies that change the distribution of resources and affect consumption inequality would also affect the aggregate pool of resources.

A third potential extension would be to recognize that Social Security benefits are typically paid to two members of a household, because both husbands and wives are eligible to receive benefits. Recent work by Gustman and Steinmeier (1999) suggests that a very substantial part of the redistribution that the current Social Security program carries out on an individual basis is undone when one considers redistribution across households. The households with higher lifetime incomes often have secondary earners who receive net benefits with larger present discounted values (in relative terms) than the secondary earners in lower-income couples. Some of this effect operates through differential mortality of secondary earners in different types of households, and some operates through differences in lifetime earning histories. The central message of this work is that considering couples rather than individuals is a key step for analyzing Social Security policies. The framework developed in this paper could be extended to consider marriage and to create households with two earnings streams.

A fourth, and particularly ambitious, extension of the current work would involve using actual earnings histories rather than simulated earnings histories to evaluate Social Security redistribution. There is undoubtedly more information in the history of actual earnings processes than in the stylized set of earnings histories generated by the stochastic models in the present chapter. Actual earnings histories are increasingly available for research purposes—for example, in conjunction with the Health and Retirement Survey. It would be intriguing to learn whether the patterns of consumption inequality that emerge in the current analysis are broadly confirmed if the analysis is based on actual earnings experience.

The authors of the present chapter have tackled one of the central issues in the analysis of Social Security reform. There is little doubt that substantive discussions of Social Security reform in the political arena will turn not only on distributional issues across generations, which have received much attention in the academic literature, but also on how reform will affect the distribution of resources within cohorts. This chapter provides a very valuable set of insights for addressing within-cohort redistribution, and it is sure to stimulate further work.

References

Gustman, Alan L., and Thomas L. Steinmeier. 1999. How effective is redistribution under the Social Security benefit formula? Dartmouth College, mimeograph.

Poterba, James, and David Wise. 1998. Individual financial decisions in retirement saving plans and the provision of resources for retirement. In *Privatizing Social Security,* ed. Martin Feldstein, 363–93. Chicago: University of Chicago Press.

Discussion Summary

Martin Feldstein asked about the relative importance of the timing of contributions, because those who contributed early would get the advantage of more years of compound interest. The authors explained that with high wage growth, most income occurs at the end of the life cycle. Consequently, almost all contributions are made at the end of the life cycle, and the impact of compound interest is negligible. The more significant factor is inequality in earnings, not timing of contributions.

John B. Shoven felt that it was important to examine inequality of opportunities, but not necessarily inequality of outcomes. If people are informed and given a broader menu, then the extra choice allows people to maximize their utility better: It should not be argued that increased choice is terrible because it increases inequality. The authors responded with two points. First, most of the inequality comes from earnings inequality, which is not really chosen in this model. Second, society seems to be concerned about choices that are made voluntarily, but are very bad decisions ex post. For instance, Teachers Insurance and Annuity Association (TIAA-CREF) has increased the choices, but has limited those choices in some ways because it is worried about people's voluntarily choosing to become very poor.

Laurence J. Kotlikoff indicated that the traditional life-cycle picture, high nonasset income followed by lower nonasset income during retirement, does not seem to correspond with data in the United States. If Social Security, Medicare, and Medicaid are part of net nonasset income, cohorts going through the life cycle now have net nonasset income that rises through retirement. The number of people in old age who are borrowing constrained seems to be important. Hence, an accurate profile of net nonasset income is vital to examining this issue. In addition, if Social Security is actuarially unfair, then the ceiling on taxable earnings implies that the lower-income classes are affected more severely by this unfairness. Very wealthy people have avoided most of the actuarial penalties of Social Security, and their consumption will be proportionately higher when compared to everyone else. In this situation, Social Security might actually increase consumption inequality. The authors acknowledged the second observation, but said it would be difficult to build an earnings ceiling into their model for various technical reasons.

Alan Gustman mentioned that when redistribution is examined by family, about half of the redistribution in the current system disappears. The amount of redistribution within the family depends on the number of years worked by the spouse. While he could understand the difficulties preventing the inclusion of these factors in the model, Gustman believed that a large amount of risk sharing and redistribution occurs between family members.