

Social Security and Inequality over the Life Cycle

Angus Deaton
Pierre-Olivier Gourinchas
Christina Paxson

Princeton University

September 1999

Prepared for NBER Conference on the Distributional Effects of Social Security Reform, Woodstock, Vermont, October 22–23, 1999. Deaton and Paxson acknowledge support from the National Institute of Aging through a grant to NBER, and from the John D. and Catherine T MacArthur Foundation within their network on inequality and poverty in broader perspectives. We are grateful to Laurence Kotlikoff, Jeffrey Liebman and Martin Feldstein for helpful comments and discussions on an earlier draft.

0. Introduction

This paper explores the consequences of social security reform for the inequality of consumption across individuals. The basic idea is that (at least part of) inequality is the consequence of individual risk. 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 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 the paper. 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 different social security arrangements on inequality among members of any given generation.

A concrete and readily analyzed example is 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, 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, though 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 for example 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; there are both theoretical and empirical studies implicating 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 1 works entirely within the framework of the

permanent income hypothesis (PIH). We derive the formulas that govern the spread of consumption inequality, and shows 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 US, but at much the same rate in Britain and Taiwan. The PIH is convenient because it permits closed-form solutions which show explicitly how social security is related to inequality. However, it is not a very realistic model of actual consumption in the US, 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 that leave the present value of lifetime resources unchanged have no effect on consumption. In consequence, in Section 2, 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 consumers endogenously switching 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.

1. Social security and inequality under the permanent income hypothesis

Section 1.0 introduces the notation and basic algebra of the permanent income hypothesis, while

Section 1.1 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 1.2 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 $(1 - \tau)^2$. Sections 1.3 discusses what happens when there is a maximum to the social security tax, and Section 1.4 extends the model to deal with finite lives and retirement and shows that the basic result is unaffected.

1.0 Preliminaries: notation and the permanent income hypothesis

It is useful to start 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

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

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:

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

for expectation operator E_t conditional on information available at time t . It is convenient to start with the infinite horizon case; the finite horizon case is dealt with in Section 1.4 below.

That consumption follows a martingale follows from manipulation of (2);

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

“Disposable” income y_t^d is defined as earnings plus income from capital

$$y_t^d = \frac{r}{1+r} A_t + y_t \quad (4)$$

Saving is the difference between disposable income and consumption

$$s_t = y_t^d - c_t \quad (5)$$

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

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

Assets are linked to saving through the identity (implied by (1) and (5))

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

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

$$\alpha(L)(y_t - \mu) = \beta(L) \epsilon_t \quad (8)$$

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 2, we will work with a process with a unit root but specified in logarithms.)

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

$$\Delta c_t = \eta_t = \frac{r}{1+r} \frac{\beta \left(\frac{1}{1+r} \right)}{\alpha \left(\frac{1}{1+r} \right)} \epsilon_t \quad (9)$$

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.

1.1 Spreading inequality

Start from the simplest illustrative case where earnings are white noise, and add an i suffix for an individual

$$y_{it} = \mu_i + \epsilon_{it} = \mu_i + w_t + z_{it} \quad (10)$$

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 i.i.d. over time. Given (10), equation (3) implies

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

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

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

so that consumption inequality is increasing over time.

Note that although (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 (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 (6),

$$s_{it} = \frac{\epsilon_{it}}{1+r} \quad (13)$$

while assets satisfy

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

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

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

which implies, after some manipulation, that

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

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

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

From (14), the variance of assets satisfies

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

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 US, 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 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 25 to about 0.38 (0.62) at age 55. We shall return to these findings in Section 2 below.

1.2 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 2.1. We recognize that the establishment of personal

accounts has other effects, some of which are not captured under our simple assumptions. But our concern here is with the reduction in the pooling or risk sharing that is implied by removing 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 accounts, but they reduce the risk-sharing elements of the current system.

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

$$y_{it} = (1 - \tau)(\mu_i + \epsilon_{it}) + \tau \bar{\mu} \quad (19)$$

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

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

Compared with the original earnings process (10), there is a shift toward the grand mean—the redistributive 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 equations (11), nor the saving equation (13), the asset equation (14), and the 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 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-section 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 40 years, the imposition of a social security tax at 12.4 percent will reduce the cross-sectional standard deviation of consumption by a factor of 3.16 (the square root of the 40 year average of 0.767 to the power of t from 0 to 40.)

In (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, (20) becomes

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

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

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

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

1.3 Social security with a maximum

The permanent income hypothesis is not well-suited to modeling a social security system where 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 gets above the social security maximum, and one whose earning never gets 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

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

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

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

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, though it will bring the two groups closer together than they would have been in the absence of the system.

1.4 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 (6). When there is a unit root in earnings, earnings immediately prior to retirement has a unit root, and so does the drop in earnings at retirement. In consequence saving, which has

to 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. But this seems more a matter of degree than an essential difference.

People work until age R and die at 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

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

From (25), we can write

$$c_t = c_0 + \sum_{s=0}^t \beta_s^{-1} \eta_t \quad (27)$$

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

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

With the social security scheme, after tax earnings while working is

$$y_{it} = (1 - \tau)(\mu_i + \epsilon_{it}) = (1 - \tau)\mu_i + (1 - \tau)(w_t + z_{it}). \quad (29)$$

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

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

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 if either (a) the autocorrelation structure of the macro component of earnings were such that current innovations had information about what will happen in retirement; this seems like an issue that is hardly worth worrying about, or (b) with precautionary motives or borrowing restrictions, where 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 2. 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” savings 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.

2. Social security with precautionary saving or borrowing constraints

2.1 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

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

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

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

where N is the number of people and Y_j is aggregate before-tax earnings for the cohort in year j .

This must equal the present value at R of social security payments, which is

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

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

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

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

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

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 reparametrize the system. Suppose that V_j , 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 the each annuity V_i to the present value of contributions gives the relationship between α and ϕ ,

$$\phi = \frac{\tilde{\alpha}}{\tau} \sum_{j=R}^T (1+r)^{R-j}. \quad (36)$$

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 that is sequestered into personal accounts. Equation (34), which constrains the parameters of the social security system, can be rewritten in terms of ϕ as

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

Note also that the individual social security payment (31) can be rewritten

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

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 does 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. And 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 $\tau(1 - \phi)$, 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.

2.2 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 (1999) and Ludvigson and Paxson (1998), 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 through $R - 1$, 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

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

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) non-stationary process

$$\ln y_t = \ln y_{t-1} + \gamma + \epsilon_t - \lambda \epsilon_{t-1} \quad (40)$$

which derives from a specification in which log income is the sum of a random walk with drift γ and white noise transitory earnings. Consumers are assumed not to be able to borrow, which requires a modification of (39), see below. One reason is to mimic the US, 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 (1998) 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 done 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 40 year working life. After retirement, there is no further uncertainty and consumption can be solved analytically for each of the 20 years remaining. We assume that the social security system presented in Section 2.1 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 (1), evolves during the working life $t < R$ according to

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

During retirement, for $t \geq R$,

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

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

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

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 $t < R$,

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

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

satisfies, for $t \leq R$,

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

and is constant thereafter. If we define σ_t , the ratio of S_t to current earnings and thus the “social security replacement rate” so that, the corresponding evolution equation is

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

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

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

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 (47) to solve backwards 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 (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.

2.3 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, $r = 0.03$, $\delta = 0.03$; (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 \$32,896 for the 90th.

Figure 1 shows the averages over the 1,000 consumers of the simulated 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 are 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 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 2 also shows more clearly than Figure 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 and/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 have poor earnings draws 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 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 results are given in numerical form in Table 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 equivalence 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 preferences, 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.

Figure 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 life-time 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. While there is no theoretical reason that prohibits this, we have not so far developed a convincing explanation of why it should occur.

Overall, the results in Figure 3 and Table 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 a quarter of social security taxes going to personal accounts, the gini coefficient of consumption increase 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 US 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 rose from 0.37 in 1981 to 0.41 in 1986.

Figure 4 compares our simulated patterns of inequality over the life-cycle with those calculated from the data in the Consumer Expenditure Survey 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 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 5 and 6 we turn to the life-cycle pattern of inequality in assets, in Figure 5 for assets excluding social security wealth, and in Figure 6 including social security wealth. The permanent income model is excluded from these comparisons since average wealth is negative for much of the life-cycle. Total wealth at any given age is defined as the sum 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 40, 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 Figure 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 diminishes with age. Even so, the most autarkic systems are the most unequal at all ages.

2.4 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 placed 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 choice or because of differential management fees, then a move to personal accounts can be expected to increase inequality by more than in the analysis so far.

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 these at the outset of their working life. The eleven mutual funds have rates of return from 2.5 to 3.5 percent a year. One can think of the funds as having identical (S&P 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 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 in the first year of retirement for the two interest rates is less than \$5,000. This difference is for a spread of a full 1 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 7 shows the case where consumers are distributed over (fixed) rates of return from 1 percent to 7 percent, compared with the case when all get 3 percent. This can be thought of as the case where 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 paper.

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.

3. List of works cited:–

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. 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.
- 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, 1999, "Consumption over the life-cycle," NBER Working Paper No.7271, Cambridge, MA. (Jul.)
- Ludvigson, Sydney and Christina H. Paxson, 1998, "Approximation bias in linearized Euler equations," NBER Technical Paper 236, Cambridge, MA. (Mar.)
- Wilkinson, Richard, 1996, *Unhealthy societies: the afflictions of inequality*, London, Routledge.

Table 1a Gini coefficients for consumption, with different social security plans
Isoelastic preferences

“Age”	G=0	G=\$5000	G=\$10,000	G=\$15,000	\$G=20,000
Isoelastic preferences, $r=.03$, $\delta=.05$, borrowing constraints					
25	0.37	0.37	0.37	0.37	0.37
29	0.41	0.40	0.40	0.40	0.40
34	0.44	0.43	0.43	0.43	0.43
39	0.46	0.46	0.46	0.45	0.45
44	0.48	0.47	0.47	0.47	0.46
49	0.50	0.50	0.49	0.49	0.48
54	0.52	0.51	0.51	0.50	0.49
59	0.52	0.51	0.50	0.49	0.48
64	0.52	0.50	0.47	0.44	0.41
69	0.52	0.50	0.47	0.44	0.41
74	0.52	0.50	0.47	0.44	0.40
79	0.52	0.50	0.46	0.43	0.40
84	0.52	0.49	0.46	0.43	0.39
Isoelastic preferences, $r=.03$, $\delta=.03$, borrowing constraints					
25	0.37	0.37	0.37	0.36	0.36
29	0.40	0.40	0.40	0.40	0.39
34	0.43	0.43	0.43	0.42	0.42
39	0.46	0.45	0.45	0.44	0.44
44	0.47	0.46	0.46	0.46	0.45
49	0.50	0.49	0.48	0.48	0.47
54	0.52	0.51	0.50	0.49	0.48
59	0.52	0.51	0.49	0.48	0.48
64	0.52	0.50	0.47	0.45	0.42
69	0.52	0.50	0.47	0.45	0.42
74	0.52	0.50	0.47	0.45	0.42
79	0.52	0.50	0.47	0.45	0.42
84	0.52	0.50	0.47	0.45	0.42

Table 1b Gini coefficients for consumption, with different social security plans
Quadratic preferences

“Age”	G=0	G=\$5,000	G=\$10,000	G=\$15,000	G=\$20,000
Quadratic preferences, $r=.03$, $\delta=.03$, borrowing constraints					
25	0.38	0.38	0.38	0.38	0.38
29	0.41	0.41	0.41	0.41	0.41
34	0.44	0.44	0.44	0.44	0.44
39	0.47	0.47	0.47	0.47	0.47
44	0.49	0.49	0.48	0.48	0.48
49	0.51	0.51	0.50	0.49	0.49
54	0.53	0.52	0.51	0.50	0.49
59	0.53	0.51	0.50	0.49	0.48
64	0.52	0.49	0.46	0.43	0.39
69	0.52	0.49	0.46	0.43	0.39
74	0.52	0.49	0.46	0.43	0.39
79	0.52	0.49	0.46	0.43	0.39
84	0.52	0.49	0.46	0.43	0.39
PIH: Quadratic Preferences, $r=.03$, $\delta=.03$, no borrowing constraints					
25	0.37	0.36	0.36	0.35	0.34
29	0.41	0.40	0.39	0.38	0.37
34	0.45	0.44	0.43	0.42	0.41
39	0.48	0.47	0.46	0.45	0.44
44	0.51	0.50	0.48	0.47	0.46
49	0.55	0.54	0.52	0.51	0.50
54	0.58	0.57	0.55	0.54	0.53
59	0.58	0.57	0.55	0.54	0.53
64	0.58	0.57	0.55	0.54	0.53
69	0.58	0.57	0.55	0.54	0.53
74	0.58	0.57	0.55	0.54	0.53
79	0.58	0.57	0.55	0.54	0.53
84	0.58	0.57	0.55	0.54	0.53

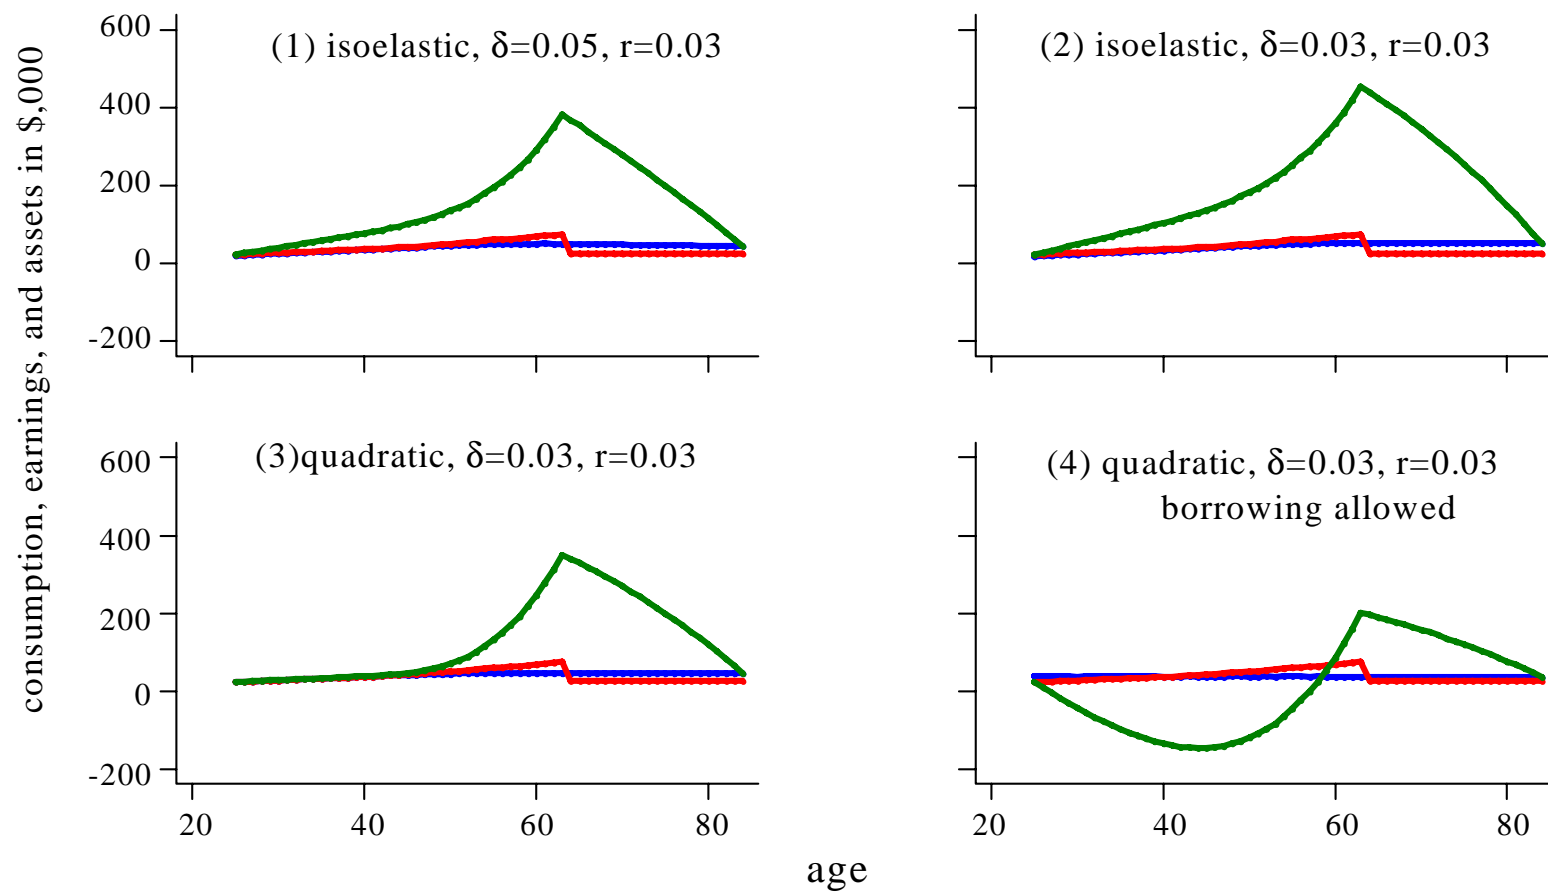


Figure 1: Age profiles of consumption, earnings, and assets for difference specifications, $G=\$5,000$

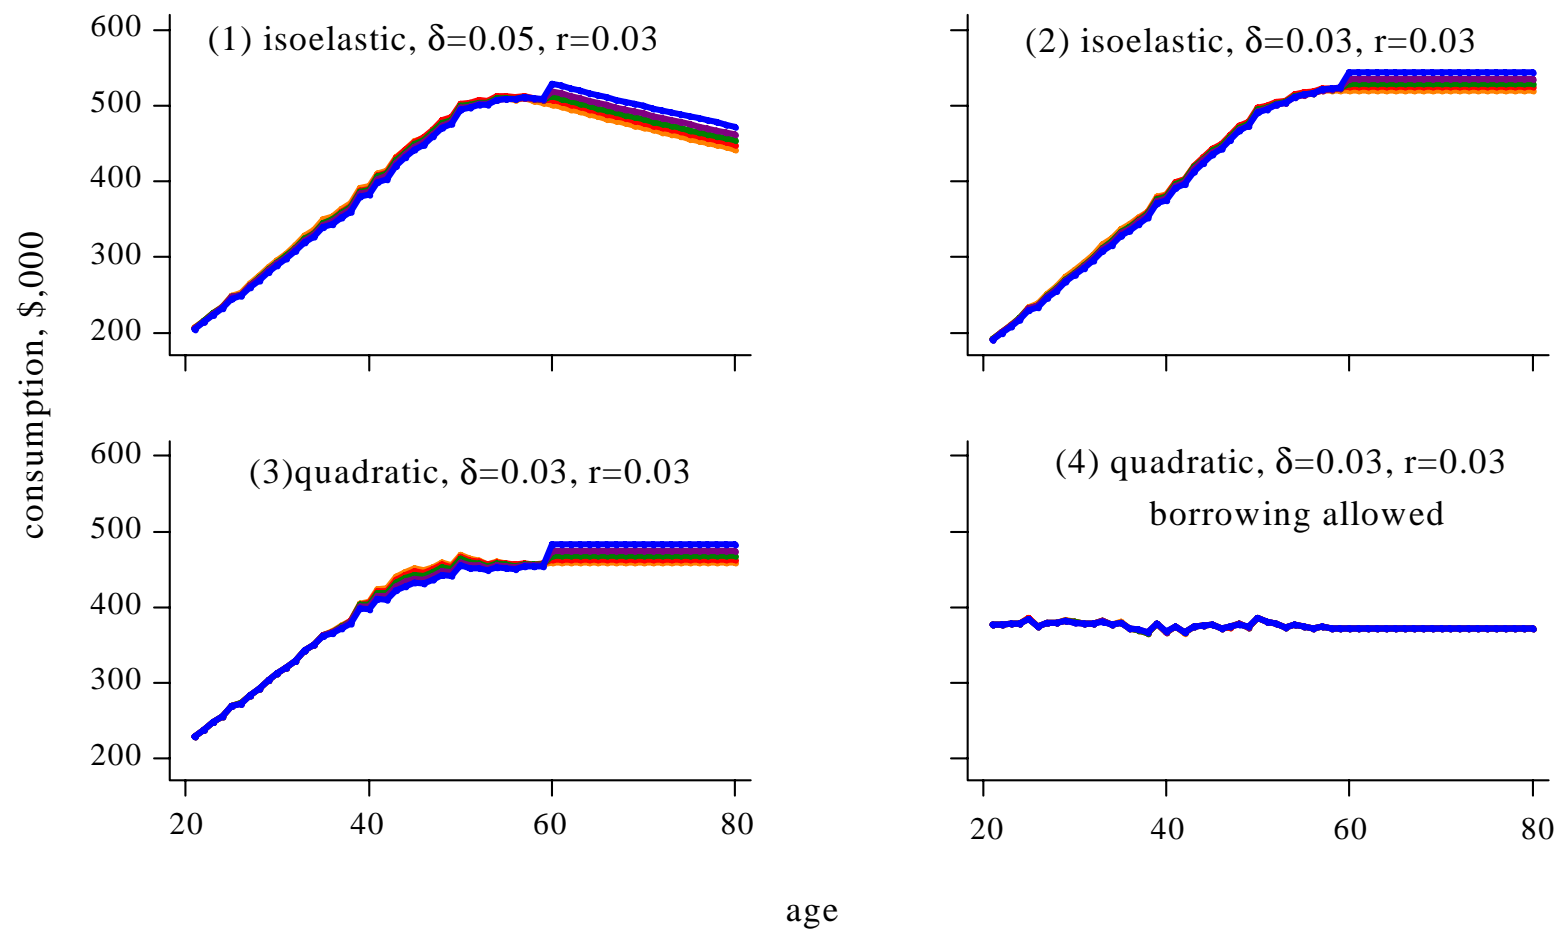


Figure 2: Consumption profiles under different specifications and alternative social security rules

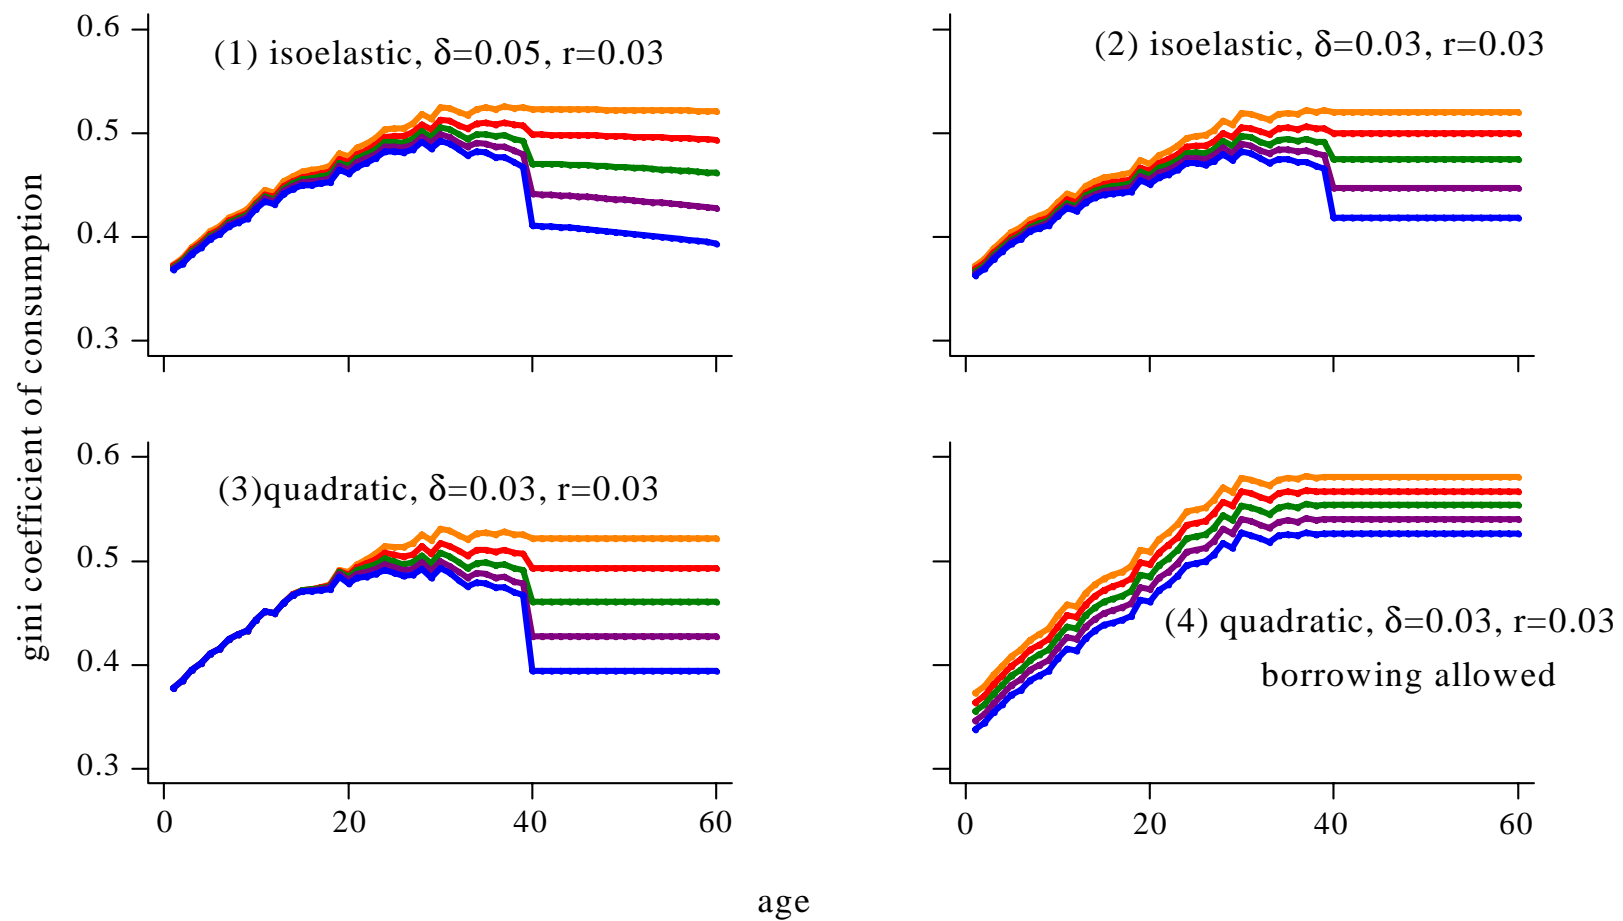


Figure 3: Consumption inequality for different specifications and social security systems

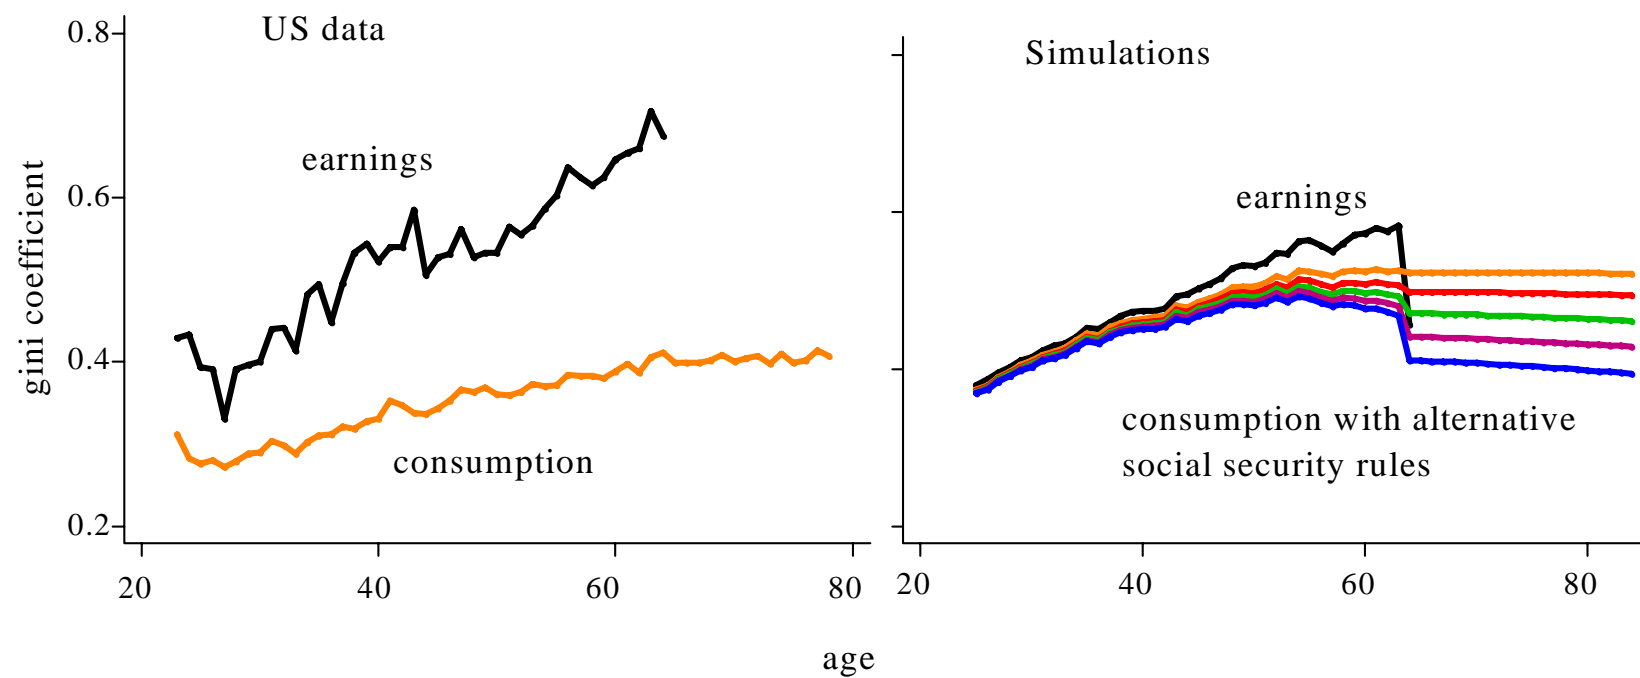


Figure 4: Actual and simulated inequality of earnings and consumption

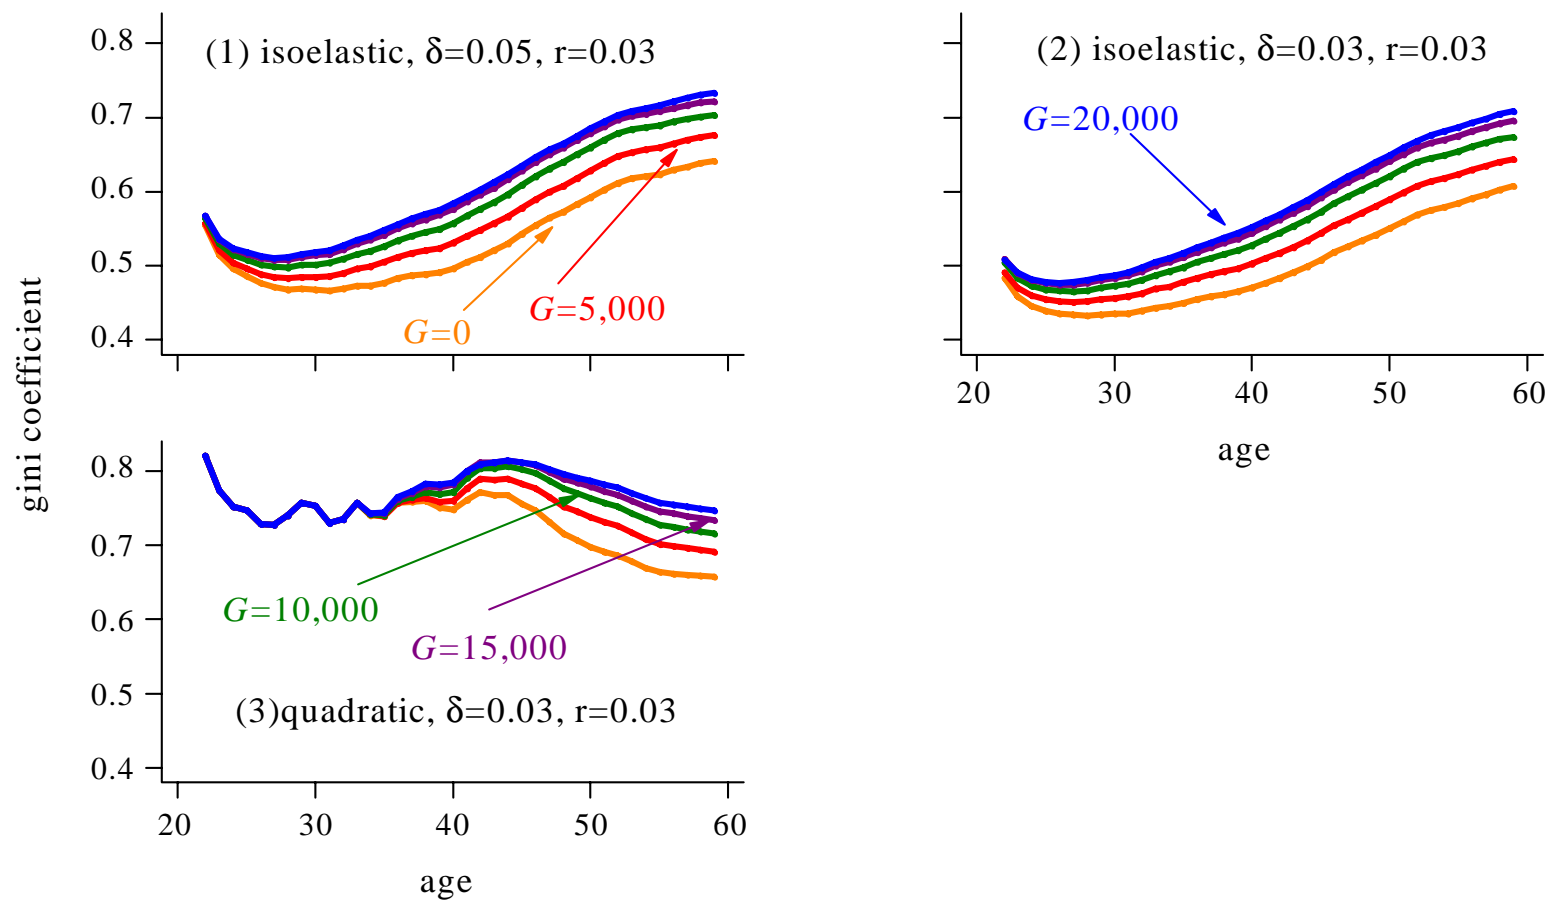


Figure 5: Gini coefficients for assets excluding social security assets

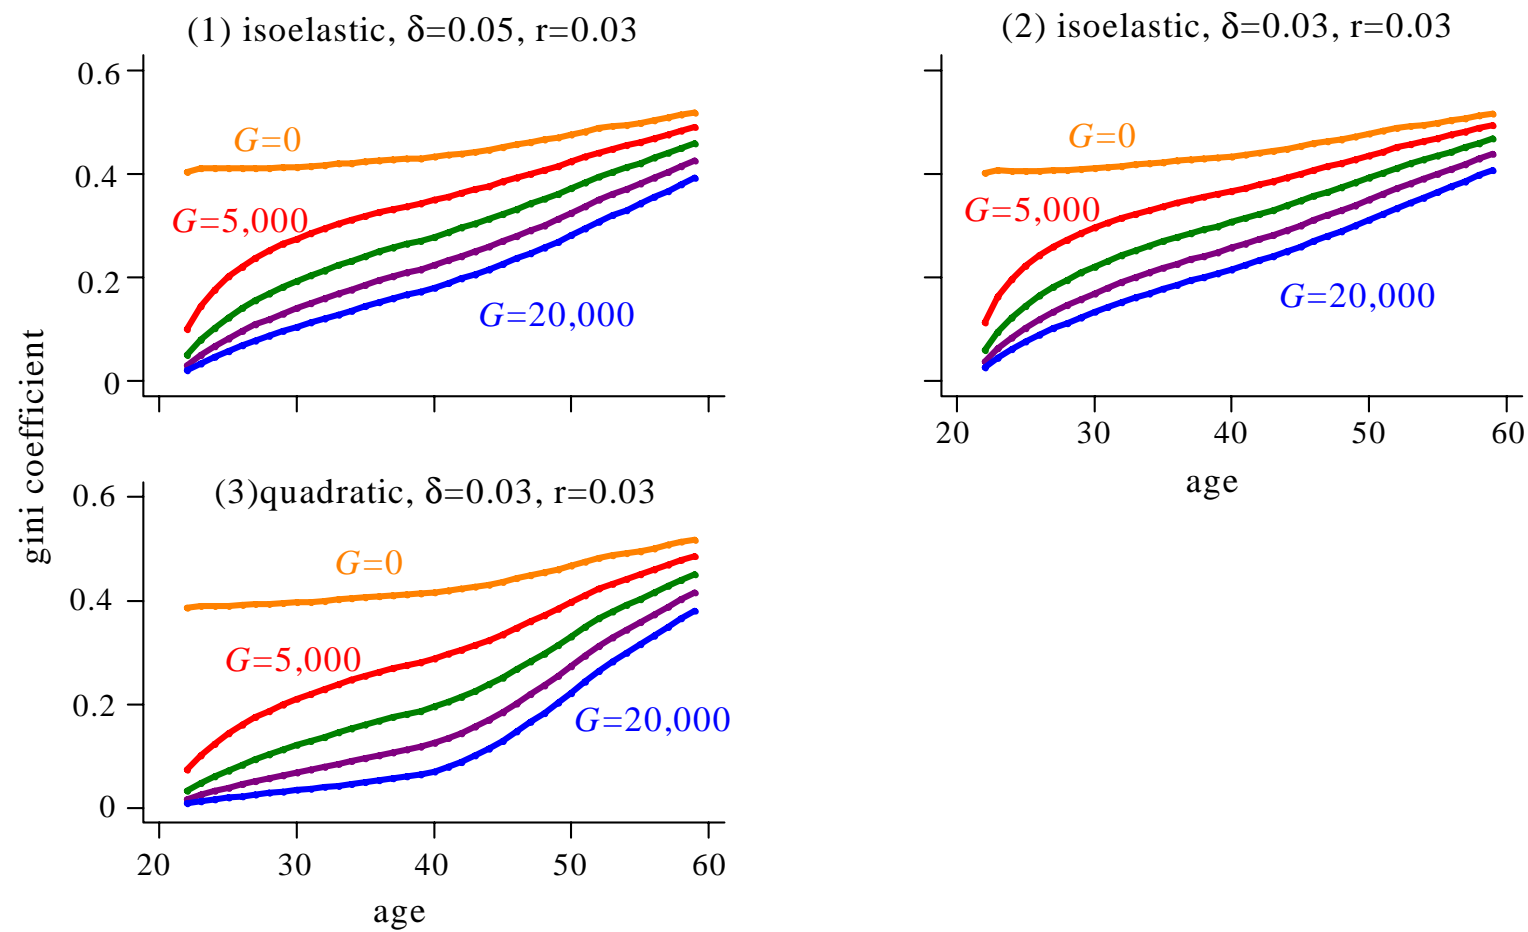


Figure 6: Gini coefficients for total assets, including social security assets

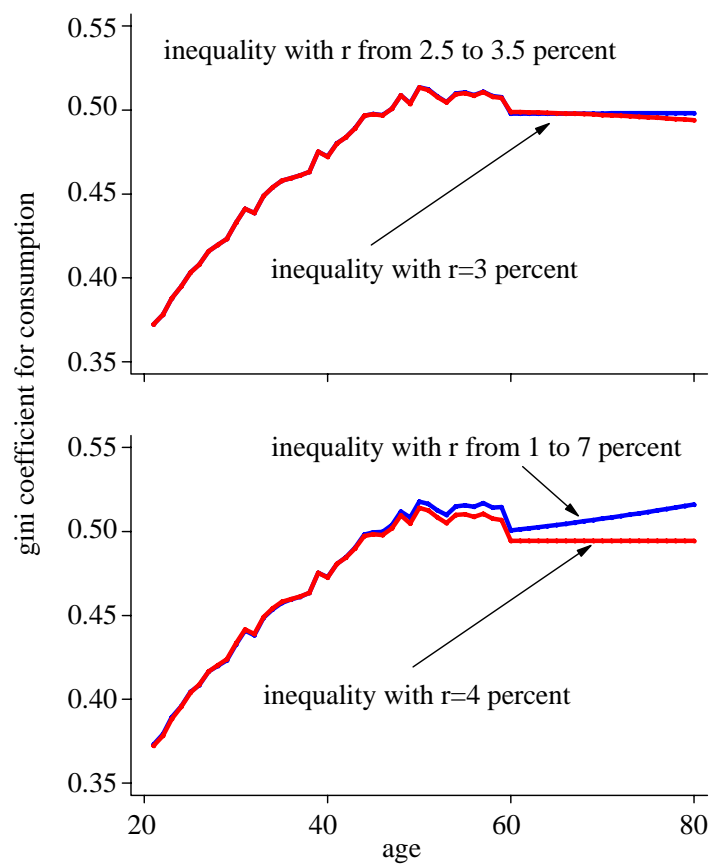


Figure 7: The effects on consumption inequality of a distribution of interest rates