Preliminary

DIFFERENTIAL MORTALITY AND THE VALUE OF SOCIAL SECURITY

David Pattison
Social Security Administration
July 25, 2002

david.h.pattison@ssa.gov

Division of Economic Research
9th Floor, ITC Building
500 E St. SW
Washington, D.C. 20254-0001

This paper was prepared for the NBER Social Security Working Day, July 31, 2002. The opinions are those of the author and should not be construed as representing the opinions or policy of SSA. The author would like to thank Dean Leimer, Joyce Manchester, Hilary Waldron, Jeffrey Liebman and Alain Jousten for helpful discussion, comments, and correspondence.

Abstract

This paper explores some issues surrounding the effect of differential mortality on the lifetime value of lifetime Social Security taxes and benefits. A small microsimulation sample of lifetime earnings histories and the associated Social Security benefits and education-differentiated survival probabilities is used to illustrate the points. Conventional money's-worth studies find that higher mortality reduces the value of lifetime benefits. In contrast, policy simulations comparing Social Security lifetime benefits to alternative arrangements do not necessarily find that higher mortality groups are at a disadvantage from Social Security. It is clear on reflection that "money's worth" is not a stand-alone concept, that any money's-worth calculation is implicitly a comparison with some alternative policy, and that if alternatives do not include annuities that perfectly differentiate by survival probabilities then neither should the money's-worth calculations. A closely related issue is the correct relation between utility value of benefits and the actuarial value. Bernheim (1987) argued that if marginal saving does not use annuities then the utility of benefits is best approximated by simple discounting without the use of survival probabilities. Bernheim's argument extends to the case in which annuities are freely available but are not perfectly differentiated, in which case utility is best approximated by actuarial values that do not use differentiated mortalities. A constant relative risk aversion (CRRA) utility specification with no borrowing constraints and no bequest motive, which is very easy to compute on a microsimulation sample, demonstrates these points. The comparison of utility changes across income classes supports the widespread use of the ratio of net transfers to lifetime income in the study of Social Security progressivity. Going beyond the simple specification, the effects of borrowing constraints and bequest motives are explored. Borrowing constraints, while they reduce the average value of lifetime net transfers, do not appear to introduce much additional variation from differential mortality. Bequest motives, on the other hand, do have differential mortality effects on utility, although a bequest motive strong enough to make the standard actuarial valuation accurate introduces some unresolved complications into the analysis.

Summary

This paper explores some issues surrounding the use of mortality differentials in the calculation of the lifetime value of Social Security taxes and benefits. The points are illustrated with a small microsimulation sample of male lifetime earnings histories and the associated social security benefits and education-differentiated survival probabilities.

Conventional money's worth studies calculate a lifetime present value of Social Security benefits by multiplying the benefit payable at each age by the probability of surviving to that age. (An additional interest rate discount is also applied.) When survival probabilities differ between two groups the lifetime values are affected. Because males, for example, have lower survival probabilities in old age than females, the lifetime value of benefits for males is lower than that for females with the same annual benefits.

Policy simulations comparing Social Security lifetime benefits to alternatives with mandatory unisex annuities do not necessarily find that higher mortality groups are at a relative disadvantage from Social Security. Because unisex annuities do not increase payments to males to compensate for their shorter expected lifetimes, the lifetime value of the annuity payments to males will be lower than that to females who have the starting principle and hence the same annuity paments. Although males appear to get a worse deal from Social Security than similarly situated females, they get the same worse deal from unisex annuities.

It is clear on reflection that "money's worth" is not a stand-alone concept and that any money's-worth calculation is implicitly a comparison of Social Security benefits with some alternative. If the alternative is an annuity that does not perfectly differentiate by survival probabilities, then neither should the money's-worth calculations. The conventional money's-worth approach is correct when the alternative is an annuity that perfectly differentiates by mortality.

Using a utility analysis, it can be shown that, in the absence of a bequest motive, the difference in lifetime utility attributable to the Social Security taxes and benefits is almost exactly proportional to the present value of the taxes and benefits, as long as that present value is constructed using the discounting that is implicit for each worker's saving for retirement. If workers would not use annuities in the absence of Social Security, then the appropriate present value uses only interest-rate discounting without adjusting for survival probabilities. If workers would use unisex or other non-perfectly-differentiated annuities for their incremental saving, then the appropriate present value uses the corresponding non-differentiated survival probabilities. Only if workers would use perfectly differentiated annuities for their incremental saving should the differential mortality survival probabilities be used in calculating the present value of Social Security benefits. If, for example, male workers in the implicit alternative would be allowed to purchase gender-specific annuities, then the money's worth calculations should include the lower male survival probabilities. These conclusions hold true even though the worker's utility takes into account the differential mortalities, as long as workers do not value any assets left behind when they die.

The calculation of the lifetime utility changes is easy to carry out on a simulation file of lifetime earnings, taxes, and benefits, allowing graphical depictions of the lifetime value by income or by mortality group. The utility changes can be depicted either directly, in which case they are sensitive to a "risk aversion" parameter in the specification of the utility function, or they can be depicted as "equivalent variations," in which the utility changes are converted into the dollar equivalents, which are much less sensitive to the risk aversion parameter. The direct utility change depiction, although it is in unfamiliar units, supports the

widespread use of proportional variations or the ratio of net transfers to lifetime earnings in the analysis of Social Security: for at least one commonly-used value of the risk aversion parameter the proportional variation replicates the utility change.

The simulation of the value of benefits in the absence of annuities and bequest motives illustrates Bernheim's (1987) claim that annuities and Social Security are best valued using simple discounting without mortality if marginal saving does not use annuities. Although Bernheim mentioned that bequest motives would invalidate this argument, the quantification of bequest motives for money's worth calculations has been unduly neglected. However, Jousten (2001) recently pointed out that a strong enough bequest motive can have the effect of inserting differential mortalities back into the present value calculation even when annuities are not used for saving.

Using an extension of Jousten's bequest motive that allows the exploration of different degrees of altruism and different average bequests at different lifetime earnings levels, I am able to illustrate Jousten's point. There are complications, however, that hinder a straightforward application of bequest motive adjustments to money's worth calculations. The question deserves further study, and we can only conclude tentatively that in the presence of bequest motives the true discounting lies somewhere between discounting without any survival probabilities and discounting with differentiated probabilities.

I. Introduction

This paper explores some of the issues surrounding the value of Social Security benefits in the presence of differential mortality. A small simulation file of male lifetime earnings and benefits is used as an exploratory tool. The only mortality differential used in the current analysis is differentials by education. The conclusions, however, should generalize to more general settings.

Two different but related problems motivated the research. The first stems from a difference between conventional money's worth studies and policy microsimulations of Social Security alternatives featuring mandatory unisex annuities. The money's-worth studies typically find that low-survival subgroups of the population receive lower aggregate lifetime transfers from Social Security, other things equal. The simulations of mandatory annuities, because the unisex annuities do not differentiate mortalities, also show lower lifetime payments to low-survival subgroups, other things equal. The money's worth calculation can be thought of as comparing Social Security benefits with an implicit alternative, and in the conventional money's-worth calculation the implicit alternative is one in which annuities are paid that are perfectly differentiated by mortality. How should the money's-worth calculation be made if the implicit alternative is a unisex annuity or if there are no annuities?

The second problem is the distinction between the expected value of the aggregate cash transfers from the perspective of the *trust funds* and the value of Social Security transfers to the *individuals* paying the taxes and receiving the benefits. The value to the *trust funds* is a simple arithmetic aggregation, aside from interest rate discounting, of the actual payments made or expected to be made. The actual or expected payments for a cohort will depend on the survival of the cohort. Hence accurate survival probabilities are important to trust fund accounting, and if mortality differentials within the cohort are correlated with the size of the payments, accurate knowledge of those differentials will give better projections of aggregate payments.

The value of the transfers to the *individual* will depend on what other resources are available to the individual. In particular, the value of benefits is higher when other retirement income is lower, and the value of the benefits therefore depends on what other provision can be made for retirement through regular saving or the purchase of annuities. It can be shown under some common utility specifications (and in the absence of bequest valuation) that savings will be adjusted, taking into account differential mortality, in such a way that the utility value to the individual of the transfers will be approximately proportional to the sum of the lifetime payments discounted for interest and adjusted for the survival probabilities used in annuities. This is the same lifetime net transfer that is made in many money's worth studies, and the utility approach thus provides a justification for the use of net lifetime transfers. The survival probabilities that should be used in this calculation, however, are not the individual's own survival probabilities but are instead whatever probabilities are implicit in the assets available for saving for retirement income. If saving is done through perfectly differentiated annuities, then the use of similarly differentiated survival probabilities would be appropriate. If saving is done through less differentiated annuities, such as the unisex annuities that are used in many employer pension plans, then the survival probabilities should be simplified accordingly. If saving outside of Social Security would not make use of annuities, then the money's worth analysis should use only interest rate discounting, without the survival probabilities.

The two problems, that of determining the best money's-worth calculation under different implicit alternatives and that of indicating the utility to the individual of the transfers, are closely related. The utility approach is more general: it not only justifies the use of net lifetime transfers in money's-worth studies, but it can also provide some guidance in determining money's worth when the implicit alternative does not include annuities, perfectly differentiated or otherwise, or when it is assumed that there is a value to accidental bequests.

The simulation sample used to explore these issues is very simple, consisting only of males born in 1930 without imputations for the unrecorded earnings above the Social Security taxable maximum, and with survival probabilities differentiated only by education. In future work I hope to extend the sample to both sexes, to couples, and to imputed earnings above the taxable maximum and to extend the mortality differentials to sex and race. The conclusions of this paper regarding the treatment of differential mortality should apply to those more extended analyses as well.

There are some important limitations in the application of the utility approach used here. The only uncertainty modeled here is the date of death. The earnings history and date of retirement for each individual is assumed to be known from the start. Hence the only value of Social Security to the individual, aside from the transfers implicit in the progressivity of the benefit formula, is the annuity value of the constant benefits in retirement, and that component disappears if annuity alternatives are available. No precautionary savings motive is modeled, so that bequests are entirely due to life-cycle saving accumulations or intended bequests. It is assumed that workers would save for retirement if Social Security is not adequate. In reality some subset of workers might be myopic in ways that are not modeled here. These workers are problematic for any valuation of Social Security benefits, since the value to them while young is different from the value to them when they are receiving the benefits.

In the utility approach used here, a simple constant relative risk aversion (CRRA) utility function is applied to each worker on the simulation file to model the savings and utility without and with Social Security taxes and benefits. If there are no constraints against borrowing against future Social Security benefits and if there is no bequest motive, the utility calculations are straightforward, and it is relatively easy to demonstrate theoretically and to illustrate on the simulation sample some basic relationships between the present value of the Social Security payments and the utility change for each individual calculated as an equivalent variation. The utility model demonstrates, and the simulation confirms, that if savings in the absence in the absence of Social Security would not be annuitized, then the equivalent variation for the Social Security transfers is best measured with interest rate discounting that does not include survival probabilities. This last finding is actually a recapitulation of an argument made by Bernheim (1987) for the value of Social Security when savings are not annuitized on the margin. The model and simulation further show that if savings in the absence of Social Security would be invested in annuities using common survival probabilities, then the equivalent variation for the Social Security transfers is best measured with a present value that also uses common survival probabilities. This can be considered a simple extension of Bernheim's argument.

There are two important qualifications to Bernheim's argument: borrowing constraints and the valuation of bequests. I deal with these two issues in the final sections of the paper. Borrowing constraints are examined by simply imposing them on the saving problem and calculating the effect on utility. The constraints reduce the average value of Social Security transfers somewhat, but they do not appear to introduce much effect attributable to mortality differentials. Bequests are more difficult to model and have potentially larger effects. Using a parameterization of a linear bequest motive, I explore the utility of Social Security transfers under different specifications of a non-zero bequest motives. The simulations demonstrate Jousten's (2001) point that under a strong enough bequest motive the actuarial valuation, using differential probabilities, becomes the most accurate. A bequest motive this strong, however, introduces several complications into the analysis. It can be shown, however, that the differential mortality effects are small when the bequest motive is not very strong or when, even under a strong bequest motive, the valuation is restricted to the effects on utility from consumption.

Section II below describes the simulation sample and presents its estimates of the actuarial present value of the net lifetime transfers under different discounting rules. Section III surveys comparable results from studies that have examined the effects of differential mortality on Social Security and discusses the typical use of mortality differentials in money's-worth studies. Section IV applies a utility model to the simulation sample and compares the equivalent variations in utility with and without the Social Security transfers to the present value of the transfers. Section V considers some of the issues in comparing utilities across subgroups or income groups. Section VI simulates the utility changes in the presence of borrowing constraints and Section VII in the presence of a bequest motive. Section VIII concludes.

II. Simulated Net Lifetime Transfers

A simulation using an actual sample of lifetime earnings histories will be used to illustrate the points in the rest of the paper. Although the sample gives a realistic distribution of lifetime earnings, and the Social Security benefits that are calculated from it make use of the current-law U.S. benefit formulas, the immediate goal is not so much a study of the U.S. benefit system as an examination of the techniques that can be used to study such a system. Hence, no great attempt has been made to exactly replicate actual transfers, particularly in this preliminary version, which only has male earnings and the resulting worker benefits. The tax rate in the simulation is set to balance the benefits in the sample under steady growth.

The simulation sample is described in more detail in the data appendix. Briefly, it consists of a sample of 298 Census males (275 after certain low earners and workers with no earnings before age 30 are removed) who reached age 62 in 1992, matched to SSA administrative data on earnings from 1951 through 1999. No calculations can be made in this sample for auxiliary (spouse or widow) benefits. Hence the current sample should be considered illustrative of worker benefits rather than representative of all benefits.

In the simulations workers can earn from age 21 through age 69 but will be assumed to consume from age 30 through age 100. Earnings before age 30 are included in lifetime wealth and build up in a fund that can begin to be consumed at age 30. Worker-financed consumption is specified to begin at age 30 for two reasons. First, it avoids dealing with the issue of how much parental and other support finances consumption before age 30. Second, for the simulation of borrowing constraints, consumption before the first year of earnings would cause problems, and the simulation file guarantees that all workers will have earnings by age 30.

The simulation focuses on steady-growth pay-as-you-go Social Security taxes and benefits under assumptions of constant price and earnings growth, constant real interest rates, and constant population growth. The following assumptions were used in the current analysis: prices grow at 3.3 percent per year, average real wages grow at 1 percent per year, and the real interest rate is 2 percent per year. The rate of steady-state population and employment growth, used to derive the steady-state tax rate, is 0.5 percent per year. The growth rate of the aggregate real wages is therefore about 1.5 percent per year. The sample earnings were adjusted slightly to replicate these steady growth assumptions. Lifetime benefits were then calculated from the adjusted earnings. The payroll tax rate was set to finance the aggregate sample benefits on a pay-as-you-go basis. (The resulting payroll tax rate was 11.7 percent. The inclusion of women and of couples benefits would have led to a higher payroll tax rate.)

Although there are interesting questions surrounding the analysis of progressivity in the startup generations of a pay-as-you-go Social Security system, or in the transition to a system with more funding, or in the steady state of a funded system, the current paper focuses entirely on the steady-state of a pay-as-you-go system. The interest rate and economic growth assumptions are such that the actuarially discounted net benefits will on average be slightly negative, although they will be positive for low earnings workers.

For the remaining analysis, all dollar amounts were converted into 1999 dollars. The earnings at age x is denoted by y_x, taxes by tax_x, benefits by ben_x, and net transfer by $t_x = ben_x - tax_x$. Because the benefit is normally zero while working and the tax is zero while retired, the net transfer is normally the negative of the tax while working and the benefit while retired.

Two sets of survival probabilities are used. One set, denoted L, gives the male cohort 1930 cohort survival probabilities, with the probability at age 21, L_{21}, set to 1.

The other set, denoted M, gives a set of survival probabilities differentiated by education, described in the appendix. Three education groups are distinguished: less than high school, high school or some college, and four or more years of college. The three different mortality rates (in log form) are shown in Figure 1A. These log mortality rates are converted into probabilities of survival to age x, M_x. The survival probabilities M_x scaled so that $M_{21} = 1$ are shown in Figure 1B.

Education is the only mortality differential currently applied; later extensions may add race and income. (Sex and possibly marital status will be added when the sample is extended to women.)

Simply-discounted and actuarially-discounted lifetime values

The real interest rate r used in developing the steady-growth scenario is also used to calculate a simple discounting series by age, R_x, declining from the value at age 21, which is set to 1:

$$R_x = 1 / [1+r]^{x-21}.$$

The value of a stream of payments y_x, discounted to age 21, is

$$Y_R = \Sigma_{x=21}^{100} R_x\, y_x\ .$$

To simplify the notation, the range of integration will usually be left out:

$$Y_R = \Sigma_x R_x\, y_x.$$

Simple discounting applied to earnings (y_x), benefits (ben_x), taxes (tax_x), and net transfers ($t_x = ben_x - tax_x$) gives the lifetime values:

$$Y_R = \Sigma_x R_x\, y_x.$$

$$Ben_R = \Sigma_x R_x\, ben_x.$$

$$Tax_R = \Sigma_x R_x\, tax_x.$$

$$T_R = Ben_R - Tax_R = \Sigma_x R_x\, t_x.$$

The interest-rate discounting series R_x also describes asset accumulation at the interest rates implicit in R:

$$1{+}r_x = R_{x-1}/R_x.$$

An asset at age x-1, k_{x-1}, will grow by age x to

$$k_x = [1{+}r_x]\, k_{x-1} = [R_{x-1}/R_x]\, k_{x-1}$$

or

$$R_x\, k_x = R_{x-1}\, k_{x-1}.$$

"Actuarialized" rates of return, which pay higher returns to survivors and nothing to the assets of the deceased, using the annuity survival probabilities L, are notated similarly:

$$k_x = [R_{x-1}/R_x]\, [L_{x-1}/L_x]\, k_{x-1}$$

$$= [R_{x-1}L_{x-1}] \,/\, [R_xL_x]\, k_{x-1}$$

$$= RL_{x-1}/RL_x \;\; k_{x-1}.$$

The notational shortcut, RL_x for R_xL_x, will be used for other similar combinations.

Discounting for common survival gives lifetime amounts that will be denoted by an RL subscript, corresponding to the earlier simply-discounted valued denoted with an R subscript:

$$Y_{RL} = \Sigma_x RL_x\, y_x.$$

$T_{RL} = \Sigma_x RL_x t_x$.

In the same way, the lifetime actuarial values using differentiated mortalities will be denoted with an *RM* subscript:

$Y_{RM} = \Sigma_x RM_x y_x$.

$T_{RM} = \Sigma_x RM_x t_x$.

A '0' superscript will be used to denote variables in the absence of Social Security transfers and a '1' superscript to denote variables in the presence of Social Security transfers. The Social Security system changes a person's lifetime wealth from $W^0 = Y$ without Social Security to $W^1 = Y + T$ with Social Security, that is, by the value of the net transfer. Progressivity studies often focus on the transfer as a proportion of initial lifetime earnings, T/Y, or $(W^1\text{-}W^0)/W^0$. If this ratio falls as Y rises, the transfer is considered progressive.

Figure 2A shows, plotted against a measure of lifetime earnings (y_{RL}, an annualized mortality-adjusted measure to be described below), the values for the lifetime taxes Tax_{RM} plotted as blue '-' signs, the lifetime benefits Ben_{RM} plotted as black '+' signs, and the net benefits, equal to the benefits points minus the tax points, plotted in three symbols, '*1*' for the lowest education group (less than high school), '*2*' for the middle group (high school) and '*3*' for the upper group (four or more years of college).[1] The tax points form three straight lines, differentiated very slightly for the three mortalities. The benefit points separately more widely with mortality because the differences between survival probabilities are cumulatively larger at older ages. Three lines representing the three mortalities can be discerned. (The additional scatter around the three lines is due to variation in the average earnings used in the benefit formula, the AIME, relative to the average lifetime earnings used on the horizontal axis.) The shape of the three benefit lines reflects the three-bracketed PIA formula, rising steeply at first relative to the tax line, then more shallowly.

The net benefits by themselves are repeated, at a larger vertical scale, in Figure 2B. The same three education groups are shown, and (color-coded) "smooths" are added. The effects of bendpoints in the benefit formula show up in the smooths, although because of a scarcity of college educated males near the first bendpoint the smooth does not show that corner well. Several other lines are added to Figure 2B and to the corresponding figures later. A colored arrow at the right side of the graph, labeled with its corresponding digit, is plotted at the average net benefit for each education group in the sample. A horizontal black dashed line indicates the overall sample average. A horizontal solid line is plotted at 0.

The horizontal axis in these and subsequent plots is a sort of annualized lifetime earnings,

---

[1]In the color version of this paper, they are also plotted in three colors, green (less than high school), yellow (high school) and red (college). When actuarial valuations separate the points into bands with college uppermost, the colors will be in traffic-light order: green on the bottom, yellow in the middle, and red on top.

$$y_{RL} = \{ \Sigma_x RL_x y_x \} / \{ \Sigma_x RL_x \}.$$

Because the common survival probabilities are the same for all, the series *RL* is the same for all, and the annualized earnings is proportional to lifetime earnings Y_{RL}. Plotting with Y_{RL} would give exactly the same picture except for a less interpretable scale on the horizontal axis. Other candidates for the horizontal axis would be Y_R, Y_{RM}, or their annualized versions. Because most earnings occur before there is a large differentiation in survival probabilities, using Y_{RM} rather than Y_{RL} would make little difference. The annualized version, y_{RM}, would reduce the high-survival earnings more than the low survival earnings, shifting the college (red) points slightly to the left in the plots. The correct value to use is unclear. It could be argued that of two persons with equal lifetime earnings but different life expectancies after retirement, the one with the longer lifetime is effectively poorer, since the resources have to be spread over a longer period, and that the correct picture would therefore be given using the differential mortality, y_{RM}. There is unlikely to be agreement about this, however, and because the visual differences are small (I have checked) I present only one version.

The position on the horizontal axis determines who is to be considered equal in a progressivity analysis. The vertical axis indicates how equal persons are treated, and the same question arises. The vertical axis, too, could be annualized, and if annualized using differential mortalities would shrink the high-survival longer-life persons toward zero lifetime benefits and a smaller net transfer. The differences (again, I have checked) do not have a noticeable effect on the separation by mortality group in plots like Figure 2B.

The lifetime transfer discounted with differential survival probabilities, as presented in Figure 2B, represents the conventional money's-worth analysis. At any given level of lifetime earnings, college educated males have a higher lifetime transfer, high school educated males have less, and high school dropouts have still less. The higher transfers to college-educated males are enough to overcome the progressivity of the benefit formula, so that they receive higher net transfers, on average, than do high school males. This is not true, in this sample anyway, of high-school males compared to high-school dropouts.

In Figure 3C, alternative measures of the lifetime net transfer are plotted. Figure 2B is repeated on a smaller scale as Figure 3C. Figure 3A is the net transfer under simple discounting, T_R. Figure 3B is the net transfer discounted for survival at the common survival probabilities, T_{RL}. All three plots have the same horizontal and vertical scale.

The simply-discounted net transfer in Figure 3A represents the analysis advocated by Bernheim (1987). If workers cannot annuitize their savings, if they are not borrowing constrained in the presence of old-age benefits, and if they do not value bequests, then the closest expenditure equivalent to the utility of their future taxes and benefits is calculated without the use of survival probabilities. This will be verified more precisely in Section IV, where the utility equivalent variations in the absence of annuities and borrowing constraints will be calculated and displayed.[2]

[2]The calculations here use a maximum lifespan of 100. The simply-discounted lifetime values are slightly sensitive to the lifespan assumption, and a maximum span of 119 gives a slightly different Figure

Comparing 3B with the simply-discounted values in 3A, the most dramatic effect of applying common survival probabilities is the overall reduction in the lifetime net benefit. Under simple discounting it is positive even at high incomes, but when survival probabilities are incorporated, depressing the value of benefits relative to taxes, all net transfer values are reduced, and at higher incomes they become negative. Just as in 3A, however, there is little or no indication of differences by education at a given income. Figure 3C, which applies the mortality differentials by education, presents a dramatic difference. Although the overall trend is the same as in Figure 3B, with negative lifetime benefits at high incomes, there is now a fairly sharp separation by education. Many or most of the college-educated males in the sample have positive net benefits. The average difference by education spans a larger proportion of the distribution of individual net benefits than in Figure 3B.

The overall effect of the differential mortality adjustment is clear to the eye even without the aid of the numerical averages indicated by the horizontal arrows. In Figure 3B the averages are close together, but in Figure 3C the college-educated have a substantially higher average. In this sample, furthermore, the college-educated no longer have a negative average lifetime transfer.

Figure 4 gives corresponding plots of the ratios of net transfer to earnings, T/Y, which is often used in progressivity studies. Again, all three plots have the same scale. The tilt is downward, indicating progressive net transfers. (This result will be standard for male worker benefits. The studies that have found severely reduced progressivity have included couples benefits and have repositioned couples with non-working wives to higher positions in the income distribution to reflect the incomes that might have been earned if the wives had worked more.) It is not possible to tell, eyeballing the data, whether the differential mortalities also decrease the progressivity of the net benefits, the downward tilt from left to right. If all of the higher educated workers had high incomes and all the lower educated workers had low incomes, the upward shift of the survival-adjusted net benefits would modify the general downward tilt. But there is too much overlap among the education groups for a change in the tilt to be detectable by eye.

Aggregating net lifetime transfers and net lifetime earnings by subgroup, the ratio of transfers to earnings using the common survival probabilities is -1.6 percent for the less than high school group, -2.3 percent for the high school group, and -1.9 percent for the college group. Using the differentiated mortalities, the corresponding percentages are -2.1 percent, -1.6 percent, and +0.2 percent. In other words, the introduction of mortality differentials causes net transfers as a percent of earnings to decline by 0.5 percentage points for the less than high school group but to increase by 0.7 percentage points for the high-school group and 2.1 percentage points for the college group. The difference between college and less than high school changes from -0.3 percentage points without the differentials to 2.5 percentage points with the differentials, an increase of 2.8 percentage points.

Dividing the sample into five quintile groups by lifetime earnings, and calculating the ratio of lifetime transfers to lifetime earnings in each group under common survival probabilities, there is a strong

---

3A, with less of a fall at higher incomes. The survival discounted lifetime values are not much affected by payments above age 100.

progression by lifetime earnings in the ratio, changing from +4.0 percent in the bottom quintile, to -0.6 percent, -1.7 percent, -2.7 percent and -3.5 percent in the second through fifth quintiles. Using differential survival probabilities, the ratios do not change much: +4.2 percent in the bottom quintile, and, in the second through fifth, -0.6 percent, -1.5 percent, -2.4 percent, and -3.1 percent. The spread between the ratios in the bottom and top quintiles narrows from 7.5 percent without differentials to 7.2 percent with differentials, a change of 0.3 percentage points, compared with 2.8 percentage points for the change in the spread by education.

The introduction of mortality differentials in this sample thus has a sizable effect on the transfers by education group but a much smaller effect on the progressivity of the transfers.

III. Money's-worth Studies, Policy Simulations, and Differential Mortality

Differential mortality effects similar to those in Figures 3C and 4C are often found with regard to other mortality differentials in progressivity studies of Social Security lifetime net benefits. Just as the mortality differentials by education have a sharply distinguishable effect on the average lifetime net benefit by education level but a less striking effect on progressivity, so do other differentials that don't explicit differentiate on income. In particular, the introduction of race differentials into survivor probabilities, while it has modest effects on the overall progressivity of net benefits, has much more notable effects on the average net benefit by race.

A number of money's-worth studies with mortality differentials have examined the results with and without the differentials, and bear out this general conclusion. Meyer and Wolff (1987) carry out analyses with three different sets of mortality tables: unisex tables, tables differentiated by sex and race, and tables differentiated by sex, race, income, education, and marital status. They compute a "redistributive component" of Social Security benefits, calculated as the portion of the benefit in excess of the annuity that would be payable from each person's accumulated lifetime taxes, using the three sets of life tables to calculate three possible annuities. Under the unisex tables, nonwhites have a substantially higher redistributive component than whites. The move from the unisex tables to the sex and race tables reduces the difference between nonwhites and whites considerably, and the move to the more completely differentiated tables reverses the difference. When classified by income, in contrast, Meyer and Wolff find so little difference among the three sets of survival tables ("almost no effect") that they present the results by income only for one set of tables.

Duggan, Gillingham, and Greenlees (1995) estimate mortality differentials by lifetime earnings on an administrative data sample and simulate net lifetime benefits without and with their estimated differentials. They find that the net transfers by income class after the introduction of their estimated lifetime earnings differentials are "noticeably less" progressive than before the introduction of the differentials but still "strongly progressive."

Coronado, Fullerton, and Glass (2000) differentiate mortality rates by income, but include many additional factors affecting progressivity. Although they find that the combination of all their factors renders the system regressive, the contribution of differential mortality itself is quite small.

Liebman (2001) differentiates mortality by race, sex, and education. Comparing the net transfers and the internal rates of return with and without differential mortality, he finds that blacks have an advantage over whites that is sharply reduced when the differential mortalities are introduced. A similar sharp reduction shows up in the higher rates of return of less-educated workers relative to higher-educated workers, and the difference in returns between high-school and more than high school disappears altogether when the differentials are introduced. The progressivity of the benefits, however, measured by the variation between lifetime earnings quintiles of the ratio of lifetime benefits to lifetime taxes, is only very slightly affected by the introduction of the differentials.[3]

Gustman and Steinmeier (2000, 2001) present two versions of the same paper, the first without mortality differentials by income and the second with differentials by income. Several of the redistributional measures show almost no change between the two papers. (There are sharp changes in at least one measure, but, given the lack of change in other measures, this must be due to some change other than to the introduction of mortality differentials.)

Cohen, Steuerle, and Carasso (2001) using a sample of administrative/projected earnings data and the benefits calculated from those earnings, examine lifetime net benefits without and with mortality differentials. (The differentials include sex, education, race/ethnicity, and a measure of current household income relative to average household income by age and marital status.) They find that the relative advantage of some race and education groups is undone by the introduction of differentiated mortality. Using internal rates of return, for example, before including mortality the return for male high-school dropout worker benefits is 3.51 percent and for college education 2.92 percent. After including mortality the returns are, respectively, 1.88 percent and 2.52 percent. The effects by lifetime income quantile are much smaller, particularly for the later of two cohorts that were analyzed, and do not come close to undoing the effects of the progressive benefit formula.

In summary, the money's-worth studies indicate that the introduction of mortality differentials can sharply reduce the measured lifetime net transfer to race or education groups with lower survival probabilities. The effect on the progressivity of lifetime net transfers relative to lifetime earnings appears to be much smaller, even in those studies that explicitly introduce mortality differentials by lifetime income. The widespread notion that mortality differentials have a large impact on progressivity might be due to several factors. First, the notion that the progressivity of the benefit formula could be reversed by longer lives of higher earners was plausible on its face and seemed to be borne out by simulations on representative workers. Second, recent studies of progressivity or redistribution of Social Security benefits have usually incorporated differential mortality, and although the strong effects on progressivity in these studies come from the treatment of single-earner married couples and of years without earnings, a hasty reader might miss the fact that differential mortality itself plays only a minor role. Third, there may be a tendency to underestimate the overlap in the earnings distribution between high and low mortality groups, so that findings of demonstrable effects on blacks or low-educated males lead to a mistaken conclusion that progressivity will be strongly affected.

[3]This measure is not explicitly calculated in Liebman, but is easily computed from his Table 4. The ratio of benefits to taxes, 1.41 in the bottom quintile and .86 in the top quintile under unisex differentials, becomes 1.39 in the bottom quintile and .87 in the top quintile with the sex, race, and education differentials.

The education differentials used in this paper are more useful for considering the effects of mortality differentials that are independent of lifetime earnings than they are for considering the effects of differentials that are strongly correlated with lifetime earnings. The education differentials cause a shift in level at each point on the lifetime earnings (horizontal) axis in Figure 3C or 4C. A differential related to earnings would, in contrast, affect the tilt of the plot. For the purposes of this paper, the shifts in level, easily coded visually into three groups, are handy for interpreting the results. That the education differentials do not, it seems, change the overall progressivity very much, is in accord with the other earnings-sample studies cited above.[4]

It should be kept in mind, however, that the education differentials used here present a very simplified picture of the possible effects of mortality differentials. Widening the scope to include race and income differentials, although it would cloud the plots with too many mortality types to easily analyze, would give a more accurate picture of the dispersion from mortality effects and of the possible effects on progressivity. Widening the scope still further to include women's benefits and the benefits of married couples would cloud the picture even more but would bring the analysis more into line with more comprehensive studies of Social Security lifetime transfers. The points that are made in this paper about education differentials, however, extend to other mortality differentials.

I argue in this paper that Figure 3B is a more correct indication than Figure 3C of the value to the individual of the Social Security transfers, unless annuities differentiating by education are available. The money's-worth studies, in giving results closer to Figure 3C, are implicitly assuming that annuitized saving vehicles exist that fully compensate low-mortality groups for their shorter life expectations by paying them higher interest rates on their savings.

Any study of the money's worth of Social Security can be considered a policy simulation comparing the taxes and benefits under Social Security with the payments under some benchmark alternative, but the alternative is rarely described explicitly. The implicit alternative can, however, be deduced from the actuarial calculations. The implicit benchmark alternative that has been used in almost all money's-worth studies, it turns out, differs from the policy alternatives that have tended to be used in policy simulations that aim at realistic descriptions of alternatives to current Social Security.

Policy simulations of Social Security alternatives have tended to feature annuitized payments in retirement. A representative example of such policy simulations is Feldstein and Liebman (2000), which compares a version of current-law benefits with alternatives based in part on annuities from investment

[4]Waldron, who supplied the education morality differentials, warns that simultaneously estimated education and lifetime earnings differentials would give different results. The education differentials would probably shrink, and the lifetime earnings differentials would give an effect on progressivity that is missing here. The use of education differentials by themselves should be considered primarily illustrative. Firm conclusions about the effects of mortality differentials on progressivity can only come from studies that either estimate the differential by lifetime earnings and use it without other factors (like Duggan et al.), or that estimate the earnings differentials simultaneously with any other differentials that are being used.

accounts. The annuities are calculated using unisex life tables. The simulated streams of current-law benefits and alternative benefits and annuities are analyzed using mortalities differentiated by sex, race, and education. (The differential mortalities are used at several places: they determine the probability of receiving widow benefits, they are used in the construction of a simulated cross-section of beneficiaries, and they are used in the tabulation of internal rates of return.) Feldstein and Liebman note of their "PRA" annuities that "although an actuarially fair PRA system would give each individual the same rate of return, we ... assume that the PRA annuities would be calculated using a single uniform unisex mortality table. The PRA system therefore gives a higher rate of return to those groups that have higher life expectancies..." In the comparisons with current-law benefits, therefore, the mortality differentials net out.

The important thing to note here is that the policy simulations quite naturally make use of both the more aggregated survival probabilities, *L*, for the calculation of the annual annuities, and the more detailed differential probabilities, *M*, for the determination of how long people live and the aggregate lifetime benefits they are likely to receive. They do it this way because that is how the probabilities occur in real life. If annuities were likely to be calculated with sex-differentiated, or sex- and race-differentiated, or sex- and race- and income-differentiated probabilities, then the simulations would seek to use the same sort of probabilities. Because annuities are not calculated this way, the simulations use the more appropriate common survival probabilities.

This is not, however, the procedure that has been followed, explicitly or implicitly, in money's-worth studies. Some money's-worth studies come close to being parallel policy simulations of the sort just described, calculating the "counterfactual" annuities that could be paid from accumulated payroll taxes and comparing the resulting annual annuities with Social Security benefits. Others use a simpler non-annualized procedure, calculating the actuarial value of lifetime benefits and comparing it to the value of lifetime taxes. The two approaches, however, are equivalent, except that the simpler procedure has the advantage of abstracting from the details of annual payments.[5]

[5]The two procedures can be summarized as follows: Suppose the accumulated value of taxes at retirement is k_{62}. The counterfactual annuity approach calculates an constant annual annuity a_M using the differentiated mortality rates M_x:

$$k_{62} = \Sigma_{x \geq 62} RM_x \, a_M$$

$$\Rightarrow a_M = k_{62} / \{ \Sigma_{x \geq 62} RM_x \}.$$

This a_M is then compared each year with the Social Security benefit ben_x:

$$\text{ratio} = ben_x / a_M.$$

The non-annualized procedure skips the counterfactual annuities, calculating the lifetime value of benefits

$$Ben_M = \{ \Sigma_{x \geq 62} RM_x \, ben_x \}$$

and compares the lifetime taxes to this amount:

In all of the money's-worth studies I am aware of, however, the mortality tables used to calculate the counterfactual annuities or, in the equivalent analysis, the lifetime net transfers, go beyond the unisex mortality annuities of Feldstein and Liebman. Sex differentiated mortalities could be defended on the grounds that sex-differentiated annuities are available privately, and thus might be appropriate to comparing Social Security with some hypothetical private alternative (even if not appropriate to evaluating Social Security against an alternative government mandated annuity). But most studies go beyond sex-differentiated mortalities, incorporating race, education, income, lifetime earnings, or marital status differentials in various combinations.

The benchmark alternative that these studies are implicitly assuming is an alternative in which annuities are provided to retirees that differentiate by whatever estimated mortality differentials are available to the analysts who produced the studies. It is conceivable that in the absence of Social Security and pension annuities thriving private markets would develop that differentiated by at least some of these characteristics. This, however, is a speculation that should be made more explicit in the presentation of money's-worth calculations. If the most likely alternative, however, is annuitized saving using sex-differentiated or unisex annuities, the money's-worth calculations should be adjusted appropriately.[6]

IV. Own Utility of Social Security Net Benefits

$$\text{ratio} = Ben_M / k_{62}.$$

When ben_x is a constant real amount, the two ratios are equivalent:

$$Ben_M / k_{62} = ben \{ \Sigma_{x\geq 62} RM_x \} / k_{62} = ben / a_M.$$

The policy simulations, in contrast, calculate an annuity based on common mortality tables:

$$a_L = k_{62} / \{ \Sigma_{x\geq 62} RL_x \}.$$

Whether the comparison is at the annual level, ben_x/a_L, or summarized using common mortalities, $\{ \Sigma_{x\geq 62} RL_x\, ben\} / \{ \Sigma_{x\geq 62} RL_x\, a_L\} = ben/a_L$, or summarized using differential mortalities, $\{ \Sigma_{x\geq 62} RM_x\, ben\} / \{ \Sigma_{x\geq 62} RM_x\, a_L\} = ben/a_L$, this comparison will be different from the money's-worth comparison if $\Sigma_{x\geq 62} RM_x$ differs from $\Sigma_{x\geq 62} RL_x$.

[6] In Meyer and Wolff (1987) a unisex annuity is considered to be the most likely counterfactual annuity, but "in order to test the effect of sex, race, and other individual characteristics on the size of the redistributive component, counterfactuals were also calculated using the two more disaggregated survivor tables." If in fact the true counterfactual is unisex, I would argue that the redistributive component does not vary by sex, race, etc., except to the extent indicated using the unisex counterfactual.

In this section, the effect of differential mortality on the lifetime value of Social Security will be considered in the context of a simple life-cycle constant-relative-risk-aversion (CRRA) utility model that is easy to apply in the simulation framework. The "value" of Social Security from this perspective will be the expenditure equivalent of the difference between lifetime utility without Social Security benefits and the lifetime utility with Social Security benefits. In this section it is assumed that there are no constraints on borrowing against retirement income and no valuation of bequests. An attempt will be made in sections VI and VII to deal with these complications.

The utility perspective brings two additional factors into play compared to an analysis that simply compares the change in the lifetime budget constraint. First, it accounts for the flexibility workers have of adjusting their own saving to take account of their expected survival probabilities. A conventional policy analysis can compare Social Security benefits with the annuity payments that could be paid from contributions equal to the same taxes, but it does not account for the possibility that workers might, in the absence of Social Security benefits, augment or offset those mandated contributions by adjusting their own saving for retirement. A conventional policy analysis can apply differential mortality rates to the Social Security benefits and to the annuity payments in the lack of Social Security, reducing the actuarial present value of old-age income under either alternative, but it does not account for how much low-survival workers might reduce their own saving for retirement, effectively increasing the utility of other old-age income under either alternative. A utility framework allows more exact simulation of these effects.

Second, the utility framework allows a deeper analysis of how the value of Social Security benefits changes across the spectrum of lifetime incomes, a question that will be taken up in section V.

The utility framework depends on the specification of a utility function that we cannot observe, and the particular family of utility functions specified here depends on a parameter, ã, that would itself be unknown even if the family of utility functions were known with certainty. Nevertheless, the CRRA family of utility functions replicates some important mechanisms presumed to underlie saving for retirement, and by examining the value of benefits over a range of values for ã, we can get more insight into the actual values of benefits than we can get by looking at the actuarial values.

Utility calculations of the value of Social Security taxes and benefits are often avoided in policy analysis for at least two reasons. First, there is considerable uncertainty over the functional forms for utility that could be considered to drive people's decisions, and second, even if the form of the utility function for individuals were known, there is a major hurdle in combining or comparing these utilities across individuals. A good case can be made, however, that these hurdles are even more of a problem for the actuarial valuation of Social Security benefits. One justification for using an annuity present value for policy analysis is that under certain conditions the annuity present value is a good indicator of the change in lifetime utility attributable to a transfer program. If these conditions fail, however, then the annuity values themselves will also be unreliable indicators of the usefulness of the benefits.

The notation for lifetime actuarial values of the preceding sections is retained here for describing the budget constraint that relates consumption at each age, c_x, to lifetime earnings:

$$\Sigma_x RL_x c_x = W_{RL} = \Sigma_x RL_x (y_x + t_x) .$$

The discounts R_x and survival probabilities L_x used here should reflect the returns to saving that are available on the market. The probabilities L_x, therefore, should only be used if annuitized saving for retirement is operative, and the probabilities should reflect the probabilities actually used in the annuitized saving. Otherwise the L_x in the budget constraint should be set to 1, or, equivalently, *RL* should be replaced by *R* wherever it occurs.

A lifetime utility function is posited,

$$U = \Sigma_x D_x M_x u(c_x),$$

$$u(c_x) = c_x^{1-ã} / [1-ã],$$

in which $u(c_x)$ is a transformation of annual consumption that reflects the diminishing marginal utility of consumption, and D_x is whatever discount is appropriate to calculating lifetime utility. The utility discount D_x is analogous to the market discount R_x but can be lower in a growing economy. (In the simulations shown here, D_x was set to 1, corresponding to a utility discount of zero. Simulations with the utility discount set to 1.5 percent, compared to a real interest rate of 2 percent, gave similar results. The most visible difference was larger but qualitatively similar results from the imposition of borrowing constraints.) The subjective survival probabilities M_x are assumed here to be the same as the differentiated survival probabilities used in Section II, although if realistic behavior were being modeled they could vary from person to person.[7] The parameter ã determines the curvature of the utility function (diminishing marginal returns) and can be considered, among other things, a "relative risk aversion" parameter, hence the name "constant relative risk aversion" (CRRA) for the utility function.[8]

[7] I won't go into this problem in greater detail here except to mention it in this footnote. It is not irrational for a person to maximize utility conditional on personal survival (i.e., with M_x=1), taking expected values over states of *that person's* world, not *the* world. The values of M_x that actually determine behavior are an empirical rather than a logical question. There are a number of reasons that the subjective values of M_x might be driven down toward the objective survival probabilities. For choices among activities with different risks, the probability of death may be strongly correlated with the probability of suffering before death. The enjoyment of future consumption may be strongly (inversely) correlated with morbidity and mortality. Perhaps most important, persons exist within a network of other people who will survive them, and there may be subtle but strong hurdles to persons' treating casually their own survival when the other people do not treat it so casually. It is possible that persons without close relatives are more likely to use an M_x close to 1 than are persons with a spouse or children. The question can't easily be settled by looking at saving behavior because the persons with close relatives might also be more likely to value accidental bequests.

[8] Utility-based Social Security analyses sometimes use ã as high as 4 or 5. A convenient set of reference points is provided by the case of portfolio choice when returns to saving have a log-normal distribution and consumption is financed entirely by saving: investors with ã=0 will choose the portfolio with the highest expected return, investors with ã=1 will choose the portfolio with the highest median return, and investors with ã=3 will choose the portfolio whose modal value (i.e., the most frequent return in the distribution) is the highest.

The utility to be maximized and the budget constraint that must be met can be combined in one equation with a lagrangian multiplier $\ddot{e}$:

$$U = \Sigma_x \, DM_x \, u(c_x) + \ddot{e} \, \{ \Sigma_x \, RL_x \, (y_x + t_x) - \Sigma_x \, RL_x \, c_x \}$$

or

$$U = \Sigma_x \, DM_x \, u(c_x) + \ddot{e} \, \{ W_{RL} - \Sigma_x \, RL_x \, c_x \}$$

Note the different roles of the differentiated probabilities M_x, which directly play a role in utility, and the actuarial probabilities L_x, which directly play a role in the budget constraint and only indirectly play a role in utility. An indication of the results of this section can be derived immediately from equation (1) under the envelope theorem, according to which the effect on lifetime utility of a small change in benefits is

$$dU/dt_x = \ddot{e} \, RL_x.$$

In other words, although the annual utilities are weighted by the personal survival probabilities and utility discount, DM_x, the utility effect of a small increase in benefits is proportional to the actuarial value of the benefits using the actuarial discount RL_x . The differential mortalities M_x play a role only through any influence on the multiplier $\ddot{e}$. For the simulations to be presented, the effects on $\ddot{e}$ appear to be small. For the analysis of progressivity, as will be shown in a later section, any effect of differential mortality on $\ddot{e}$ will be swamped by the effect of lifetime income on $\ddot{e}$.

It will be assumed that the interest rates underlying R_x will not change between the economy without and the economy with Social Security (perhaps not a very plausible assumption when transfers are being made for all persons and all generations).

The above formulation gives for the marginal utility of consumption

$$u'(c_x) = RL_x/DM_x \, \ddot{e}.$$

RL_x/DM_x designates a marginal utility profile that is the same for everyone with the same M and D. The multiplier $\ddot{e}$ varies according to individual wealth and the curvature of the utility function.

This expression for the marginal utility, valid for more general utility functions than CRRA utility, gives an approximation for lifetime utility of net transfers that is easy to calculate. If we let c^0 denote the optimum consumption in the absence of Social Security transfers and c^1 the optimum in the presence of the transfers, then the lifetime utilities without and with transfers are

$$U^0 = \Sigma_x \, DM_x \, u(c^0{}_x).$$

$$U^1 = \Sigma_x \, DM_x \, u(c^1{}_x).$$

The difference is

$$ÄU = U^1 - U^0 = \Sigma_x\ DM_x\ \{\ u(c^1{}_x) - u(c^0{}_x)\ \}.$$

Using the approximation

$$u(c^1{}_t) - u(c^0{}_t) \approx (\ c^1{}_t - c^0{}_t\ )\ u'(c^*{}_t)$$

we get

$$ÄU \approx \Sigma_x\ DM_x\ (\ c^1{}_t - c^0{}_t\ )\ u'(c^*{}_t)$$

$$\approx \Sigma_x\ DM_x\ (c^1{}_x - c^0{}_x)\ ë\ \ RL_x/DM_x$$

$$\approx ë\ \Sigma_x\ RL_x\ (c^1{}_x - c^0{}_x)$$

$$\approx ë\ \{\ (Y_{RL}+T_{RL}) - (\ Y_{RL})\ \}$$

$$\approx ë\ T_{RL}.$$

Thus, under any utility function, the utility difference should be approximately proportional to the present value of the net transfer itself, discounted using the available annuity rates. The CRRA utility function, however, allows the utility difference to be calculated quickly and exactly, including the proportionality term $ë$, which might itself differ by mortality.

Under the CRRA specification the marginal utility is

$$u'(c_x) = c_x^{-ã}.$$

It follows that

$$c_x = ë^{-1/ã}\ (RL_x/DM_x)^{-1/ã}.$$

It is then easy to calculate $ë$ from the wealth constraint:

$$W_{RL} = \Sigma_x\ RL_x\ c_x$$

$$= \Sigma_x\ RL_x\ ë^{-1/ã}\ (RL_x/DM_x)^{-1/ã}.$$

$$ë^{-1/ã} = W_{RL}\ /\ [\ \Sigma_x\ RL_x(RL_x/DM_x)^{-1/ã}].$$

$$c^*{}_t = W_{RL}\ \ (RL_x/DM_x)^{-1/ã}\ /\ [\ \Sigma_x\ RL_x(RL_x/DM_x)^{-1/ã}].$$

The optimum consumption is therefore proportional to lifetime wealth, W_{RL}, and follows a path that is a function at each age of RL_x and the ratio RL_x/DM_x.

The optimum path, c^*, can be calculated both without Social Security transfers ($W_{RL} = Y_{RL}$) and with Social Security transfers ($W_{RL} = Y_{RL} + T_{RL}$). In each case, the lifetime utility can also be calculated, using

$$U = \Sigma_x \, DM_x \, (c_x^*)^{1-ã}/[1-ã].$$

Under CRRA there is a simple relationship between lifetime wealth and lifetime utility:

$$U = K \, W_{RL}{}^{1-ã}/[1-ã]$$

where *K* is a constant (for each mortality group) that depends only on *R*, *L*, *D*, *M*, and ã. This relationship is easily inverted,

$$W_{RL} = ([1-ã] \, U/K)^{1/[1-ã]}$$

allowing the quick calculation of expenditure equivalents, the starting wealth that would allow a person to reach a specified level of utility. The expenditure equivalents to the pre-transfer utility U^0 and the post-transfer utility U^1 are

$$E^0 = ([1-ã] \, U^0 /K)^{1/[1-ã]}.$$

$$E^1 = ([1-ã] \, U^1 /K)^{1/[1-ã]}$$

$$= E^0 \, [U^1/U^0]^{1/[1-ã]}.$$

The equivalent variation is

$$E^1\text{-}E^0 = \{ \; U^1 / K)^{1/[1-ã]} - \; (U^0 / K)^{1/[1-ã]} \}$$

$$= E^0 \, ([U^1/U^0]^{1/[1-ã]} - 1).$$

In the simulations the utilities are actually calculated from the optimal consumption paths and converted into equivalent variations in this way, as a check on the computations. But the values (when there are no borrowing constraints and no bequest motive) are actually much simpler to calculate, as can be seen from plugging the value for *U* as a function of *W* into the equivalent variation calculations:

$$E^0 = W_{RL}{}^0 = Y_{RL}.$$

$$E^1 = W_{RL}{}^1 = Y_{RL} + T_{RL}.$$

$$E^1 - E^0 = T_{RL}.$$

More important than the ease of computation are two implications. First, the equivalent variations are unaffected by the utility curvature parameter ã. Second, the equivalent variations are exactly equal to the net transfers calculated in Section II

Figures 5A and 5B show the equivalent variations in the utility (using γ=1.01, although it does not affect the variations) calculated under the assumption that annuities are not available (5A) and under the assumption that they are available at common life table rates L_x (5B). (Figures 5C and 5D will be referred to later.) Figure 5A clearly corresponds closely to Figure 3A and 5B to 3B. The correspondence of the A figures for equivalent variations to the A figures for net financial transfers bears out Bernheim's analysis: if annuities are not available, the utility change is accurately indicated by the discounted cash value without using survival probabilities. The correspondence of the B figures bears out the extension in this paper: when annuities are available, the utility change is accurately indicated by the actuarial value that uses the survival probabilities provided for in the annuity tables, not the actual survival probabilities faced by each person.

Although Figure 5B is identical with Figure 3B, which discounted lifetime transfers using the common survival probabilities L, it is important to note that the utilities underlying Figure 5B are discounted using the *differential* survival probabilities M. The probabilities L play a role only in determining the returns to annuitized saving. Although the differentiated probabilities M are present in the lifetime utility function, we do not see the separation in lifetime values by education that was so visible in Figure 3C.

Another version of the equivalent variation that is useful is the equivalent *proportional* variation, that multiple k of original wealth that gives the new utility:

$$U^0 = K\, W_{RL}^{1-\gamma}\ .$$

$$U^1 = (kW_{RL})^{1-\gamma}$$

$$= k^{1-\gamma}\, U^0.$$

$$k = (\, U^1/U^0)^{\,1/[1-\gamma]}.$$

Again, the actual values for the proportional variations in the unconstrained non-bequest case are a simple function of the Section II values:

$$k = T_{RL} \,/\, Y_{RL}.$$

These proportional variations are plotted in Figures 6A and 6B for the no annuities and the common-survival annuities cases. They are identical with the net transfers as a proportion of lifetime earnings shown in Figures 4A and 4B. (Figures 6C and 6D will be considered in section VI.)

Figures 5A and 6A do not show the sharp separation by education group that was seen in Figures 3C and 4C, even though the mortality differentials that caused the separation seen in those figures are fully present in the utility optimizations underlying Figures 5 and 6. To get equivalent variations like those seen in Figure 3C, the differential mortalities M_x would have to be used not just in the definition of lifetime utility but also in the lifetime budget constraint,

$$\Sigma_x R_x M_x c_x = \Sigma_x R_x M_x y_x = W_{RM},$$

implying the presence of perfectly differentiated annuities in all of retirement saving.

V. Comparing across individuals

The analysis by individual equivalent expenditure variations does not allow comparisons between persons or between subgroups classified by income or education, although the graphical presentation invites such comparison and the calculation of an average equivalent variation implicitly assumes that the variations are comparable.

The data presented in figures like Figure 5 and 6 should not, from one common standpoint, be used for much more than analyzing the sign of the utility change. A positive variation indicates a positive utility change, and a negative variation indicates a negative utility change. The equivalent variation in Figure 5B and the proportional equivalent variation in Figure 6B tell exactly the same story: the utility change from lifetime Social Security, if discounted for mortality, tends to be positive at low incomes and negative at high incomes.[9]

Strictly, the comparisons in the earlier sections are most applicable to the comparison for a given individual of the value of the Social Security transfers, holding survival probabilities constant. When comparing two persons with different survival probabilities, even if they have the same lifetime earnings, the problem of comparing different lifetimes arises. Should we focus on average utility while alive, or on total lifetime utility? The total lifetime utility is the quantity referred to in the preceding section:

$$U = \Sigma_x DM_x \, u(c_x).$$

By the "average while alive" I mean a sort of annualized utility:

$$U_{avg} = \{ \Sigma_x DM_x \, u(c_x) \} / \{ \Sigma_x DM_x \}.$$

For the individual, the question is irrelevant as long as choices do not themselves affect the length of life. For the policy maker or voter adding up utilities across persons, however, there is no clear answer. I have opted to use the annualized utility in plots of the changes in utility. Plots based on the total lifetime utility show a slight broadening in the ribbon of points in plots like 7A, exhibiting a slight mortality differential effect, although nothing on the scale of Figure 3C. I use the annualized form because it strengthens the parallelism between the net transfers as a percent of lifetime earnings (Figure 4A or 4B), proportional utility equivalents (Figure 6A or 6B), and the utility changes at ã near 1 (Figure 7A or 7B).

[9] Strictly speaking, the equivalent variations (but not the proportional variations) can play a somewhat larger role in welfare analysis. If some of the equivalent variations are positive and some negative, and if the sum across persons is positive, there is a possibility of further adjustment in the transfers so that all persons can have positive utility changes. The practical significance of this possibility might be questioned.

As ã approaches 1, the utility change in its annualized form approaches the proportional variation, which in turn corresponds to the net transfer as a percent of lifetime earnings (see Appendix for details). One implication is that analyses that use proportional variations are in effect adjusting for different lengths of life.

We now look at the utility changes themselves, rather than the equivalent variations. Because the equivalent variations are an index of the utility changes, the signs of the changes will be no different. The magnitudes, however, will be different, and as the utility curvature parameter ã is varied, the relative utility changes ÄU(ã) for high and low incomes will be altered, as is not the case for equivalent variations. The sum or average utility change will also be affected by the parameter ã.

Figures 7A and 7B show the earlier variations in terms of utility changes at ã near 1. It can be seen from comparison with Figure 6 that the equivalent proportional variations present an almost identical picture, aside from the quantities on the vertical axis.

(Plotting by total lifetime utility instead of the annualized utility, not shown here, gives, as already mentioned, a slight separation by education group in the plot corresponding to Figure 7A, due, presumably, to effects operating through the lagrangian multiplier ë. The separation is much smaller than that found in Figure 4C, and is no longer visible in the plot corresponding to Figure 7B.)

Figure 8 gives the utility changes for the unconstrained annuitization case (corresponding to Figure 7B) at ã = 1.2, ã=1.4, ã=2, and ã=3. Utility changes at higher ã are increasingly dominated by utility changes at very low incomes.[10]

If equivalent variations were calculated for ã = 1.2, ã=1.4, ã=2, and ã=3, analogous to the utility change calculations of Figure 8, they would all be identical with the plot for net transfers in Figure 3B and of equivalent variations at ã near 1 in Figure 5B. The equivalent proportional variations would be identical with the proportional net transfer of Figure 4B and the equivalent proportional variation of Figure 6B. This is because under unconstrained maximization with only cash sources of utility the equivalent variation is simply equal to the cash variation. The approximate equality between utility change and equivalent proportional variation at ã near 1, visible from comparing Figures 6 and 7, thus no longer holds at higher ã. (Looking ahead to the constrained cases considered in the next section, if equivalent variations were plotted for the constrained cases, there would be slight differences as ã increased, but nothing like the utility change differences of Figure 8.)

The lack of change in the equivalent variations as ã is increased should not be taken as meaning that the

[10]This poses more than just presentational problems. Some of the low earners in the administrative data sample really just have low *covered* earnings and might have had other earnings not covered under Social Security. The presence to an unknown degree of these non-covered earners prevents any solid conclusions about average utility changes, particularly at high curvature parameters like ã=3 or ã=5. As more recent administrative data that includes noncovered earnings accumulates, this problem will go away, but related problems regarding truly part-career earners will remain. These problems are also present for actuarial valuations, but high-curvature utility functions make them more prominent.

utility curvature parameter ã has little effect on the valuation of lifetime Social Security benefits. The equivalent variations are the utilities transformed into a dollar metric that varies for each person. The metric tells us whether the utility changes were positive or negative, but it does not tells us how the utility change for a person with lifetime income of 100,000 compares with the utility change for a person with a lifetime income of 500,000. As ã rises above 1, the equivalent proportional variation becomes increasingly inaccurate as a guide to utility changes across the income distribution.[11]

Because the proportional variation in Figure 4 is almost identical with the ã=1 utility change of Figure 7, conventional progressivity analysis can be considered a study of utility changes under the assumption that ã = 1. This amount of curvature was at one time considered perhaps the most plausible specification of utility ("Bernoullian utility"), and the advocacy of proportional and later progressive taxes may have had some of its roots in the fact that if ã=1 the burden of a tax (if measured as the utility change) will be borne equally at all income levels only if the tax is proportional to income. In more recent times a ã greater than 1 has come to be considered more plausible, and a proportional tax has come in some circles to have a regressive flavor to it, consistent with the notion that the burden of a tax should not fall more heavily on the poor, as would a proportional tax if ã were greater than 1. (See Young, 1990, for a discussion of equal sacrifice principles and an attempt to measure the utility curvature assumption implicit in income tax schedules if they are aiming at equal sacrifice.)

The fact that for net benefits a policy that appears, at ã=1, to give equal utility changes to all earners is actually, at any higher ã, a policy that gives a greater utility increase to low earners than to high earners is worth bearing in mind with regard to the money's-worth studies of Social Security. Several of these studies find that as various effects are taken account of with regard to couples benefits and the benefits of part career earners, then the net lifetime benefits come closer to being proportional to lifetime earnings. If a net lifetime transfer is proportional or barely progressive by conventional measures, this means that the utility changes are approximately constant if measured at ã = 1. If, however, the true value of ã is greater than 1, then even exactly proportional net transfers will benefit low earners more than high earners. The transfer system can in fact be somewhat regressive and still benefit low earners.[12]

---

[11]Comparisons of equivalent variations or equivalent proportional variations as the risk aversion parameter changes are for this reason very muted indications of changes in utility. In Brown (2001), for example, an "annuity equivalent wealth" measure is calculated that amounts to an equivalent proportional variation. This proportional variation is computed for different mortality groups at CRRA parameters ranging from ã=1 to ã=5. It is not surprising that the equivalent variations do not change much as ã varies. For similar reasons, one should be careful interpreting results by income class that are presented as equivalent proportional variations. Statements like the following can be found in the literature: "Utility, measured in wealth equivalents, would rise by 8 percent for average earners, 6 percent for the poorest agents, and 4.4 percent for the richest agents." The measure being discussed here is an equivalent proportional variation, and under some CRRA parameterizations, a 6 percent gain in the bottom class is a larger utility gain than an 8 percent gain in a middle income class.

[12]This remark does not apply when net transfers are negative at all income levels, in which case the rules for taxes apply: to insure that the poor do not suffer a larger utility change than the rich, the utility change must be equal at some ã *greater* than or equal to the true ã. In Coronado et al. the use of a

To some extent, the comparing of utility changes across persons that is carried out in analyses like that of Figures 7 and 8 can be considered an analysis of the expected utility changes faced by persons whose future earnings are unknown. For a young worker, Social Security provides insurance against a long career of low earnings. (This insurance component might be underrated by analysts in government and academia, whose lifetime earnings are probably more predictable than those of the average worker.) This component is completely ignored in the individual points plotted here: life earnings for each person are assumed to have been known from the start. An analysis that incorporated earnings uncertainty would be considerably more complicated and probably not feasible yet for use on administrative data earnings histories. The effect, however, must be something like the averaging together of all the points in a window around the expected lifetime earnings. As the age at which the insurance value is calculated recedes to and before birth, the window grows to include more points (and more education groups), approaching, if the process is taken far enough, the average utility change for the whole sample. The curvature of the utility function affects the average utility change in each window: the higher ã is and the more curved the utility function is, the greater will be the average utility change in each window realtive to the value in the middle of the window.

VI. Borrowing constraints

Bernheim (1987) recognized that borrowing constraints could depress the value of lifetime benefits below the simply-discounted value, and he tried to quantify how important that effect was.

The borrowing constraints that are of importance in this context are not the constraints on young earners borrowing against future income (or borrowing to invest more heavily in equities). They are, instead, the constraint on borrowing late in life against Social Security or annuity payments at the end of life, whether borrowing just before retirement against benefits after retirement or borrowing early in retirement against high benefits late in retirement. If low survival probabilities in old age lead workers to plan for high marginal utilities and low consumption in old age, the resulting optimum consumption path, if it is declining, will not sit well with a constant real benefit payment. Retirees, not allowed to borrow against these too-high payments to increase earlier consumption, are forced to a somewhat less than optimal path if Social Security or under any other mandatory annuitization sets the level or growth of benefits too high. This reduces the value of the Social Security relative to the unconstrained saving path.

This sort of constraint is relatively easy to implement in the simple saving models we have been considering. Given an unconstrained consumption path, the asset accumulation by the end of each age, k_x, can be calculated:

$$k_x = [RL_{x-1}/RL_x]\ k_{x-1} + y_x\ [\ + t_x\ ] - c_x.$$

This asset path can be checked to see if it ever drops below zero. If it does, the consumption path can be reoptimized with assets constrained not to fall below zero. This is done both without and with the Social Security transfers t_x, yielding, when the utility function is applied to the two consumption paths,

---

high discount rate made transfers negative at all income levels.

two utilities U^0 and U^1, one or both of which might be below the corresponding utility that would be calculated in the unconstrained case. It is possible for the constraint not to bind both without and with Social Security, in which case the simulated utilities are exactly the same as before, as are the calculated equivalent variations. If the borrowing constraint does bind, it is likely to either bind with Social Security but not without, or to bind more with Social Security than without, reducing the utility difference U^1-U^0.

Exact equivalent variations, conditional on the age at which the variations are supplied, could be calculated for these utility pairs. Instead, an approximation is used here. The exact pair of utilities under the constraint is calculated, but for the expenditures E^0 and E^1 corresponding to these utilities the indirect utility function from the unconstrained case is used. The difference between E^1 and E^0 is a good indicator of the utility difference, and should be close to an exactly calculated variation, although I have not checked this.

Subfigures C and D in Figures 5, 6, and 7 show the results for equivalent variations, equivalent proportional variations, and utility differences when the borrowing constraint is applied under regular saving with no annuities available (C) or under freely annuitizable saving (D). Freely annuitizable here means that annuities can be purchased one year at a time, receiving the annuity rate RL_x if the person survives, without being constrained to, for example, a constant real annuity. The plots C and D for the borrowing-constrained case have been placed to the right of the corresponding non-constrained cases, at the same scale, to facilitate comparison.

In both cases, the borrowing constraints do lower the overall equivalent for Social Security. The basic shapes, remain, however, and there seems to be little additional difference by differential mortality. Bernheim's conclusion, namely, that actuarial discounting is much less accurate than simple discounting, still stands up in the amended version, namely, that discounting at the rates used to calculate annuities is more accurate than discounting using differential mortalities.

VII. Bequests

Bernheim (1987) mentioned that bequests would invalidate simple discounting but did not deal further with that issue. Jousten (2001) analyzed the effect of bequests, finding that a bequest motive would move the correct discounting from simple discounting toward actuarial discounting, and that a strong enough bequest motive would make actuarial discounting accurate.

For simplicity, I deal in the following bequest valuation simulations only with a world in which annuitized saving is not available or is never considered an attractive alternative to leaving accidental bequests.[13] The wealth constraint is W_R, computed using the market return discounts R. All asset accumulations are bequeathable. Under the utility specification that has been used so far, these assets left behind at death had no value. It is plausible, however, that most people do not downrate accidental bequests so

[13] It should be possible to model annuitized saving vehicles that compete, through their higher returns, with bequeathable but lower-return regular savings. Full modeling of the question, however, should include the contribution to utility of balances held as precautionary savings.

severely, since the assets will be of use to their children or other recipients, and that they have in mind, when they build up and then draw down their assets, that their own consumption reduces the consumption of their heirs.

The bequest motive specification used here adds to the lifetime utility at each age *x* a function of the bequeathable assets at that age, k_x, weighted by the probability of death occurring that year:

$$U = \Sigma_x \{ D_x M_x \, u(c_x) + R_x \, [M_x - M_{x+1}] \text{ ê ë* } k_x \}.$$

The easiest factor in the added term is the probability of death during the period. If the chance of surviving to the beginning of age *x* is M_x, and the chance of surviving to the end of age *x* is M_{x+1}, then the chance of dying during age *x* is M_x-M_{x+1}.

The remaining factors could be summarized as a bequest motive linear in k_x and the probability of k_x's being a bequest:

$$U = \Sigma_x \{ D_x M_x \, u(c_x) + \text{â}_x \, [M_x - M_{x+1}] \, k_x \}.$$

This would be, at least superficially, a "joy-of-giving" bequest motive, in which the utility of the bequest is determined entirely by the amount given and not by consideration of what it might be worth to the recipient. (Hurd (1989) and Jousten (2001) also use bequest motives linear in k.) The view taken here, however, will be that there is a deeper structure to â$_x$: it is a reduced form parameter that can incorporate such considerations as the usefulness of the bequest to its recipient and how much the bequestor takes this usefulness into account. (Abel and Warshawsky, 1988, apply a similar view.) The coefficient â$_x$ will be split into three components:

$$\text{â}_x \, k_x = R_x \text{ ê ë* } k_x.$$

This formulation reflects an assumption that the value of a bequest *from* a person at a given age bears some relation to the value of a gift *to* a person of the same age. A $1 gift to a person at age *x* would increase that person's utility by R_x ë, where ë is the dollars-to-utility lagrangian multiplier for that person, which in the CRRA specification is a function of the person's wealth and the CRRA parameter ã. A gift of \$$k_x$ to the person would be worth about R_x ë k_x, especially if k_x is small relative to the person's lifetime wealth.

The value of a gift of k_x *from* the person is by analogy assumed to be R_x ê ë* k_x, where ë* is the dollars-to-utility multiplier for the presumed recipient and ê is a step-down factor equal to 1 for a perfectly altruistic bequestor who does not adjust for, among other things, the greater wealth of recipients of later generations. (Absence of a bequest motive is modeled by setting ê to 0.) The presumed multiplier ë* could be set equal to the bequestor's own multiplier ë, but more realistically it is parameterized to lie somewhere between the bequestor's multiplier and the population average ë$_{avg}$. A ë* close to the population average has the plausible consequence that high earners leave more accidental bequests than low earners.

In the simulations presented here, ë* and ë$_{avg}$ are calculated at the pre-transfer wealths of the population and do not change when transfers are included. The simulations also do not factor in the bequests received by each person. The latter should be included even when there is no bequest motive, since as long as some assets are not annuitized there will be accidental bequests in the starting endowment of each generation. These accidental bequests, however, are random variables, and do not fit well in the deterministic framework of this analysis.

Other refinements could include factoring in the possibility that recipients might not be able to fully anticipate bequests on average, or that recipients are borrowing constrained, both of which would reduce the utility of the bequest when it did occur.

The interest rate discount R_x was also used by Jousten (2001). Jousten (1999) goes into this question in greater detail in the context of modeling both gifts and bequests. A discount rate higher than the market rate would lead to making transfers as early as possible. A discount rate lower than the market rate would lead to as much postponement of transfers as possible, extending to the use of trust funds to make the transfers after death. A discount rate equal to the market rate makes the timing of the transfers irrelevant.

The largest questions surround the size of the bequest motive parameter ê. Although everyone might agree that it is nice to leave some assets behind, it is not clear how valuable they are relative to the old-age consumption that has to be foregone to leave larger bequests. Some *prima facie* evidence against a large value to bequests is the fact that the U.S. Social Security system pays only a minimal bequest amount, an amount which has been allowed to dwindle in real value over time. The value of a larger lump-sum benefit (a bequest of sorts), apparently, is felt to be small relative to the reduction in surviving worker or widow benefits that would be needed to pay the costs of the increased death benefit.

If the bequest motive is large enough, savers will plan on leaving positive assets at the maximum attainable age (here, age 100). This brings a new wrinkle to the computation of optimal consumption paths, since for savers with positive terminal assets the total lifetime consumption is less than, rather than equal to, lifetime wealth. The optimal consumption path in these cases turns out to be independent of the bequestor's own wealth (except to the extent that the ë* term depends on the bequestor's own wealth). Optimizing requires checking this unconstrained path and, if the terminal assets are negative, iteratively recomputing the path until the terminal assets are 0, an iteration that is quite feasible on the simulation sample.[14]

[14]For lifetime utility

$$U = \Sigma_x DM_x u(c_x) + \Sigma_x R_x [M_x - M_{x+1}] \text{ êë* } k_x + R_T \text{ì}_T k_T,$$

the marginal utility of consumption, using $dk_x/dc_t = - R_t/R_x$ for $x \geq t$, is

$$dU/dc_t = DM_t u'(c_t) - \Sigma_{x \geq t} R_x [M_x - M_{x+1}] \text{ êë* } R_t/R_x$$

$$= DM_t u'(c_t) - \text{ êë* } R_t M_t$$

$$\Rightarrow u'(c_t) = R/DM_t \text{ êë* } M_t.$$

In those cases in which terminal assets are positive, the only effect of a change in Social Security transfers is to change the bequests: the consumption path is unaffected. For these savers, the lifetime Social Security transfer simply adds to the expected bequest if the transfer is positive or subtracts from it if the transfer is negative. The correct survival probability for determining the expected transfer is the differentiated survival probability M_x: the longer a person actually lives, the more Social Security benefits will be accumulated and left to the heirs. The Social Security transfers therefore affect lifetime utility directly through the bequest motive terms rather than indirectly through the lifetime budget constraint and consumption. As Jousten pointed out, the actual survival probabilities enter the calculation in this way even when savers do not annuitize.

If the bequest motive is not quite strong enough to cause positive terminal capital, the lifetime utility will still partake of this effect. Some of any change in lifetime utility will be due to the change in utility from consumption operating through the budget constraint, but some will still be due to changes in the accumulation of accidental bequests and will depend on actual survival probabilities.

The utility for four bequest motive simulations are displayed in Figure 9.[15] A small bequest motive (ê = 0.2) with ë* equal to the bequestor's ë is shown in Figure 11A. The other three cases are for large bequests motives (ê = 1) at the bequestor's own ë (9B), at a ë* halfway between each bequestor's and the population average (9C) and at the population average (9D). Comparing these to the

---

If the bequest motive is weak enough that the terminal capital is zero, there will be an additional component of the marginal utility trajectory attributable to the multiplier on terminal capital:

$$u'(c_t) = R_t/DM_t\ (\text{êë}^*\ M_t + \text{ì}_T\ ).$$

But for a strong bequest motive, terminal capital becomes positive, ì_T goes to zero, and the marginal utility trajectory simplifies to

$$u'(c_t) = RM_t/DM_t\ \text{êë}^* = R_t/D_t\ \text{êë}^*$$

which is independent of wealth and mortality unless they affect ë*.

Compare this to maximization without bequest valuation but with perfectly differentiated annuities:

$$U = \Sigma_x\ DM_x\ u(c_x) + \text{ë}\ \Sigma_x\ RM_x\ (y_x - c_x)$$

$$\Rightarrow u'(c_t) = RM_t/DM_t\ \text{ë} = R_t/D_t\ \text{ë}\ ,$$

which is optimized with positive ë and k_T=0. If in the bequest valuation problem êë* were exactly at the level where k_T=0, the two problems would give êë*=ë and identical consumptions and utilities.

[15]I haven't worked out the computations yet for equivalent variations with a bequest motive.

corresponding non-bequest-motive plot (Figure 7A, on a different scale) it is apparent that a large bequest motive lowers the value of the Social Security transfers, making them negative at higher incomes. The bequest motive also introduces clear differential mortality effects, particularly when (case D) the recipient ë* is assumed equal to the population's, which causes wealthy savers to place a high value on bequests relative to their own consumption. The similarity between Figure 9B and Figure 4C is striking, and confirms Jousten's theoretical finding that for a sufficiently strong bequest motive the actuarial valuation gives a correct indication of the utility. Figure 9A, with a weak bequest motive, is much closer to Figure 4A, discounting without survival probabilities.

There are a number of features of a strong bequest motive like that behind Figures 9B through 9D that should give pause, however. First, a strong enough bequest motive, if widespread, could lead to a sort of bequest bubble, in which an increasingly large bequest is passed on from generation to generation until some generation finally decides to consume it at a much-reduced marginal utility. One of the peculiarities can be seen by imagining the introduction of a pay-as-you-go Social Security system when a strong bequest motive is widespread. Initially, a new tax is levied on the young and given to the old as benefits. Because the strong bequest motive is operative, no one's consumption changes. The old, given more benefits, simply add them to their assets, increasing their bequests. The young pay the additional taxes but do not reduce their consumption, lowering their assets at first but replenishing them later, at least on average, from their newly augmented bequests. The young will, in turn, later receive larger benefits, paid through their children's larger taxes but returned to their children through larger bequests. Although no one's consumption changes, everyone's utility increases, because under the bequest motive the larger bequests generate higher utilities. This makes for an unusual justification for pay-as-you-go Social Security.

The strong dependence of the Figure 9 results on the bequest term can be seen from Figure 10, in which only the change in utility from consumption is plotted. Other formulations of an altruistic bequest motive are possible in which the utility taken into account by the bequestor depends on the amount by which the heirs can be expected to actually increase their consumption as a result of the bequest. The parameter ë*, in other words, should really only reflect the change in the heirs' consumption utility from each dollar of bequest. Any tendency to increase bequests at the expense of consumption would be self-limiting, since a bequest bubble would be associated with a ë* of 0.

Figure 10 also reminds us that there is a sort of double counting in the imputation of the bequest motive utility to the bequestor. Lifetime utility with a bequest motive is intended to explain saving behavior more accurately. That the saver is acting to maximize this utility, however, does not mean that the lifetime utility being maximized is entirely the saver's *own* utility. Implicit in the notion of altruism is that the utility being maximized include's other people's utility, and a summation of changes in utility over a steady state population, or a plot of the changes, needs to disentangle the utilities actually achieved by members of the population from the utilities aimed at on their behalf by other members of the population.

A closely related question arises with regard to equivalent variations. When utilities are nested or interdependent, shouldn't equivalent variations be calculated simultaneously as a vector of variations, one for each member of the population in each generation? If a vector of expenditures E^1 is that vector that allows each member to attain utility U^1, taking into account the expenditure that each other member is given, then the equivalent variations, even in the presence of an altruistic bequest motive, should move back toward an indication of the changes in each member's own utility from consumption.

There are, in sum, many fascinating questions to be resolved before we can begin to state with confidence what the effects of a bequest motive are on the valuation of Social Security benefits. Bequests do provide a mechanism by which mortality differentials might enter into the proper valuation of benefits even in the absence of perfectly differentiated annuities. But whether the effect is negligible, as in Figure 9A, or large, as in Figure 9B, is, for once, not just an empirical question, but one that requires further sorting out of the theoretical questions.

VIII. Conclusion

For the calculation of Social Security money's worth, it is clear that some attention must be given to the implicit alternative to Social Security, and that the effects of differential mortality will depend on the alternative chosen. Conventional money's worth calculations typically use all the mortality differentials that can be applied, and therefore implicitly assume that annuities with the same differentials would be available in the alternative. If less differentiated annuities would be available, such as unisex annuities, then a utility analysis suggests that the unisex mortalities should be used in the money's worth calculations, even when people take more accurate mortalities into account in their consumption planning. By the same argument, if annuities would not be available in the alternative, then the money's worth calculations should use simple discounting, without adjustment for mortality.

All three varieties of discounting are easy to compute on simulation files with actual earnings and benefit data, even without recourse to utility modeling. A utility model of optimal saving and consumption, however, brings additional capabilities. It can be used to modify the money's worth calculations to take into account less than optimal annuitization payments when savers are constrained from borrowing against overly high annuities. Exploratory simulations of the value of Social Security under such a constraint, however, indicate that while the overall value is lowered slightly, mortality differentials do not play much of a role.

A utility model that includes a value for bequests shows much greater scope for introducing differential mortality effects. A model of this form, implemented on a simulation file, shows a differential mortality effect much like that of the conventional money's worth calculations that assume perfectly differentiated annuities, even when the implicit alternative assumes no annuities. Achieving this effect, however, requires a strong parameterization that might not be realistic. Weaker parameterizations of the bequest motive show results much like that for borrowing constraints: the overall value of Social Security can be reduced, but the differential mortality effects themselves appear weak. There are several open questions, however, about the strength and valuation of bequest motives that prevent any firm conclusions.

References

Abel, Andrew B. and Mark Warshawsky. "Specification of the Joy of Giving: Insights from Altruism." Review of Economics and Statistics vol. 70. February 1988.

Bernheim, R. Douglas. "The Economic Effects of Social Security: Toward a Reconciliation of Theory and Measurement." Journal of Public Economics 33,3: 273-304. August 1987.

Brown, Jeffrey R. "Redistribution and Insurance: Mandatory Annuitization with Mortality Heterogeneity." Center for Retirement Research Paper CRR WP 2001-02. http://www.bc.edu/crr. April 2001.

Coronado, Julia Lynn, Don Fullerton, and Thomas Glass. "The Progressivity of Social Security." National Bureau of Economic Research Working Paper 7520. February 2000.

Duggan, James E., Robert Gillingham, and John S. Greenlees. "Progressive Returns to Social Security? An Answer from Social Security Records." Research Paper No. 9501, Office of Economic Policy, U.S. Department of the Treasury.

Feldstein, Martin and Jeffrey Liebman. "The Distributional Effects of an Investment-Based Social Security Reform." National Bureau of Economic Research Working Paper 7492. Sept. 2000.

Gustman, Alan L. and Thomas L. Steinmeier 2000. "How Effective is Redistribution Under the Social Security Benefit Formula?" National Bureau of Economic Research Working Paper 7597. March 2000.

Gustman, Alan L. and Thomas L. Steinmeier 2001. "How Effective is Redistribution Under the Social Security Benefit Formula?" Journal of Public Economics vol. 82: 1-28.

Hurd, Michael D. "Mortality Risk and Bequests." Econometrica 57, 4: 779-813. July 1989.

Jousten, Alain. "On the Timing of Gifts and Bequests." April 1999.

Jousten, Alain. "Life-cycle modeling of bequests and their impact on annuity valuation." Journal of Public Economics vol. 79: 149-177. 2001.

Liebman, Jeffrey B. "Redistribution in the Current U.S. Social Security System." National Bureau of Economic Research Working Paper 8625. December 2001.

Meyer, Charles W. and Nancy L. Wolff. "Intercohort and Intracohort Redistribution under Old Age Insurance: the 1962-1972 Retirement Cohorts." Public Finance Quarterly, 15, 3: 259-281. July 1987.

Waldron, Hilary. "Links Between Early Retirement and Mortality." Office of Research, Evaluation and Statistics Working Paper Nb. 93. August 2001.

Young, H. Peyton. "Progressive Taxation and Equal Sacrifice." American Economic Review, 80, 1: 253-266. March 1990.

Appendix: Data and Methods

Simulation sample

The current sample consists of males in the 1993 and 1994 Census SIPP survey who reached age 62 in 1992 and who have successful matches to administrative data records on earnings. There were 301 males in the starting sample. Three had no earnings and were deleted, leaving 298 males with some earnings between ages 21 and 69. Of the 298, 9 had too few years of earnings to be insured for Social Security. Their taxes were included in the steady-state payroll tax calculation but they were excluded from the utility analysis. Another 14 had earnings starting after age 30. Their taxes and benefits were included in the steady-state calculation but they were excluded form the utility analysis because they would have caused difficulties with borrowing constrained solutions.

Earnings

The sample earnings were nominal covered earnings up to each year's taxable maximum. These earnings were first converted to relative wages by dividing by the national average wage indexing series for each year. They were then reconverted into steady-growth nominal wages by multiplying by an artificial national average wage series that had the same 1999 value as the actual average wage series but grew before 1999 at the assumed steady-state nominal per-capita wage growth (the product of the assumed real wage growth and the assumed steady price growth).

An artificial price index and real discount index were also calculated, using assumed steady inflation and interest rate assumptions. Both were set to 1.0 in 1999.

For the simulations here, the steady growth assumptions were 3.3 percent per year for prices, 1 percent per year for real wage growth, and 2 percent per year for the real interest rate. The rate of steady-state population and employment growth, used as described below to derive the steady-state tax rate, was 0.5 percent per year. The growth rate of the aggregate real wages was therefore about 1.5 percent per year.

Benefits

Earnings before age 60 were indexed to age 60 using the artificial nominal wage growth series. The highest 35 years of indexed earnings (including any unindexed nominal earnings after age 60) were averaged together and divided by 12 to get an "average indexed monthly earnings" (AIME), which was the basis for the primary insurance amount (PIA) in the benefit formula. A benefit formula close to that in U.S. law was used. Two bendpoints were calculated equal to the monthly equivalent of .33 and 1.33 of the nominal wage index in the year the cohort turned 60. The PIA was equal to 90 percent of AIME up to the first bendpoint, plus 32 percent of any AIME between the first and second bendpoints, plus 15 percent of any AIME above the second bendpoint.

The lifetime history of benefits was then calculated from the PIA. A simplifying assumption of retirement at 62 was used [in the current exercise] even for workers with earnings after age 62. The benefit at each age, accordingly, was the PIA, reduced by 20 percent for early retirement, and price indexed for the growth in prices between the year the worker turned 62 and the year of the benefit. These nominal benefits were calculated for each age through 100.

Mortality

Two sets of survival probabilities were used: the set of life-table survival probabilities *L* used to determine the rate of return to annuities and the set of mortality-differentiated personal survival probabilities *M* used to adjust individual lifetime utility for survival probabilities.

The annuity-table survival probabilities L_x were derived from the mortality rates for the 1930 male birth cohort in the 2001 Trustees Report intermediate assumptions. [This needs some tinkering. A unisex survival probability should be used, and will be used when women are included in the analysis.]

Differential mortality

The personal survival probabilities M_x were derived from differential mortalities that use the SIPP education variable for each worker in conjunction with differential mortality by education estimated by Hilary Waldron (unpublished and provisional, but see Waldron 2001 for related work, including a description of the sample and of similar logit regressions).

Waldron used the 1973 Exact Match file linked to administrative data on deaths through 1997. Each sample person provided pooled observations from 1973 through the year of death or 1997, whichever came earlier. The regressions were logit regressions on a dependent variable of 0 for each pooled year in which the person survived and 1 for the year in which the person died. The independent variables included the year of birth, dummies for three education levels (less than high school, high school degree or less than four years of college, and four or more years of college), and a series of linear age splines from age 30 to 40, 40, to 50, and so on up to age 90. The three age profiles from the regression on males, calculated for the 1930 male birth cohort, are shown in Figures 1A and 1B. Figure 1A gives the log of the estimated mortality hazard, and displays a convergence toward equal mortality at age 90 that is sometimes found in other differential mortality estimates. The highest of the three mortality profiles shown is that for less than high school education; the lowest is that for four or more years of college. For the simulation, the three paths of mortality under age 30 and over age 90 (the bounds of the mortality rates estimated by Waldron) were extended down to age 21 and up to age 100 using the growth in the corresponding Trustees Report mortality at those ages for the 1930 birth cohort.

Differentiated survival probabilities M_x were calculated from the differentiated mortality hazards q_x. The age 21 survival probability was set to 1 and the probabilities for ages 22 through 100 were calculated from that base:

$$M_x = M_{21} \, Đ_{21}^{x-1} \, (1 - q_x).$$

The resulting survival probabilities are shown in Figure 1B. For survival probabilities the profiles reverse position, the highest survival probabilities are for four or more years of college. The survival probabilities shown in Figure 1B are the probabilities M_x used in the simulations.

The summed mortality rates, $\Sigma_x M_x$, are 43.02 for less than high school, 46.03 for high school or some college, and 50.07 for 4 or more years of college. The summed interest-discounted rates, $\Sigma_x RM_x$, for the same three groups are 31.69, 32.90, and 34.52.

[In later work, these differentiated survival probabilities will be scaled so that survival probabilities averaged over the education groups in the sample at each age and sex match the overall survival probabilities by age and sex L_x.]

Steady-state pay-as-you-go taxes

To calculate the payroll tax necessary to sustain the benefits in a pay-as-you-go system, the aggregate covered earnings and benefits in a representative year are needed. The tax rate is the ratio of aggregate benefits in the representative year to aggregate covered earnings in the representative year.

The sample life histories are for a 1930 birth cohort, and show the nominal earnings and benefits for that cohort. The age 21 amounts, for example, give the nominal values for the 1930 birth cohort in 1951, the age 69 amounts give the nominal amounts for the 1930 birth cohort in 1999, and the age 99 amounts give the nominal amounts for the 1930 birth cohort in 2029. If the representative year is 1999, we need the age 21 amounts for the 1978 birth cohort and the age 99 amounts for the 1900 birth cohort.

If per-capita real earnings are growing at a steady rate *g* and prices are growing at rate *p*, the 1978 birth cohort at age 21 (in 1999) will have nominal earnings that are $[(1+g)(1+p)]^{1978-1930}$ higher than those the 1930 cohort had in 1951. If the steady-state population is growing at the rate *n*, there will be $[1+n]^{1978-1930}$ more workers in the 1978 cohort than there were in the 1930 cohort. The age 21 earnings amounts in the 1930 sample, therefore, were scaled up by a total of $[(1+g(1+n))(1+p)]^{1978-1930}$. Similar scaling was for earnings at all ages (including earnings, not in this exercise, after 1999). Finally, the earnings at each age was adjusted for the probability of having survived to that age, using the differentiated probabilities M_x. Here the differentiation is quite appropriate, since an aggregate cash value, not a welfare value, is being calculated. [In further work a slight adjustment will be made reflecting that fact that later generations have higher survival probabilities than earlier generations. This cannot be done exactly, since increasing longevity nudges the system out of a proportional growth framework, but a close approximation to the effect of increasing longevity on a pay-as-you-go tax rate should be easy to achieve.]

If nominal earnings are growing at rate g+p, the AIME for each succeeding cohort will also grow at that rate, scaling up the ensuing benefits proportionately. Although a cohort's benefits after age 62 rise only at the inflation rate *p* rather than the nominal growth rate *p+g*, and therefore fall relative to the earnings of current workers, the benefits of one cohort relative to the benefits at the same age for the preceding cohort still rise at the rate *g*. The scaling of the 1930 cohort nominal benefits was therefore exactly the same as for earnings: earlier benefits were adjusted upward by [(1+g)(1+n)(1+p)] relative to later benefits. Survival probabilities were incorporated in the same way.

Three discounting series are calculated at the beginning of the simulation, Vg, Vp, and Vn, all starting at a value in 1951 higher than 1.0 and shrinking down toward 1.0 in 1999 through successive multiplications by the factors 1/[1+g], 1/[1+p], and 1/[1+n]. The 1999 aggregates for the simulation sample are then

$$AggEarn = \Sigma_n \Sigma_x M_{n,x} * Vg(1930+x) * Vp(1930+x) * Vn(1930+x) * y_{n,x}.$$

$$AggBen = \Sigma_n \Sigma_x M_{n,x} * Vg(1930+x) * Vp(1930+x) * Vn(1930+x) * b_{n,x}.$$

Here $y_{n,x}$ and $b_{n,x}$ are the nominal earnings and benefits at age *x* for person *n* in the sample. The steady-state pay-as-you-go tax rate is

$$taxrate = AggBen / AggEarn.$$

For the current sample of 1930 males with worker benefits only, the steady-state rate is 11.7 percent at (g=.01, n=0.005, p = 0.033). A higher growth rate would give a lower tax rate. The inclusion of women in the sample together with the calculation of couples benefits would raise the necessary tax rate.

Regular or annuitized utility maximization

For no borrowing constraints and no bequest valuation, the utility function augmented with a multiplier on the constraints is

$$U = \Sigma_x DM_x u(c_x) + ë \{ \Sigma_x RL_x y_x [+ t_x] - \Sigma_x RL_x c_x)$$

$$= \Sigma_x DM_x u(c_x) + ë \{ W_{RL} - \Sigma_x RL_x c_x)$$

where the t_x is the net Social Security transfer. For the non-annuitized saving problem, replace RL everywhere (here and below) by R. For perfectly-differentiated annuities, set L equal to M everywhere.

The solution uses

$$u'(c_x*) = c_x^{-ã} = RL_x/DM_x ë_{RL}.$$

$$c_x* = ë_{RL}^{-1/ã} (RL_x/DM_x)^{-1/ã}.$$

$$W_{RL} = \Sigma_x RL_x c_x* = ë_{RL}^{-1/ã} \Sigma_x RL_x (RL_x/DM_x)^{-1/ã}.$$

$$ë_{RL}^{-1/ã} = W_{RL} / \{ \Sigma_x RL_x (RL_x/DM_x)^{-1/ã} \}.$$

$$ë_{RL} = W_{RL}^{-ã} \{ \Sigma_x RL_x (RL_x/DM_x)^{-1/ã} \}^{ã}.$$

$$c_x* = W_{RL} (RL_x/DM_x)^{-1/ã} / \{ \Sigma_x RL_x (RL_x/DM_x)^{-1/ã} \}.$$

Equivalent variations

$$U_{RL} = \Sigma_x DM_x c_x^{1-\tilde{a}} / [1-\tilde{a}]$$

$$= \Sigma_x DM_x c_x^{-\tilde{a}} c_x / [1-\tilde{a}]$$

$$= \Sigma_x DM_x [RL_x/DM_x \ddot{e}_{RL}] c_x / [1-\tilde{a}]$$

$$= \ddot{e}_{RL} \Sigma_x RL_x c_x / [1-\tilde{a}]$$

$$= \ddot{e}_{RL} / [1-\tilde{a}] W_{RL}$$

$$= W_{RL}^{1-\tilde{a}} \{ \Sigma_x RL_x (RL_x/DM_x)^{-1/\tilde{a}} \}^{\tilde{a}} / [1-\tilde{a}]$$

$$= K_{RL} W_{RL}^{1-\tilde{a}} / [1-\tilde{a}].$$

This is easily inverted into the expenditure needed to attain U_{RL}:

$$E_{RL} = ([1-\tilde{a}] U_{RL} / K_{RL})^{1/[1-\tilde{a}]}.$$

If U^0 and U^1 are the utilities without and with the transfers, then the expenditure equivalents are (dropping, for convenience, the RL subscripts on E, U, K, T, Y, and W)

$$E^0 = ([1-\tilde{a}] U^0 / K)^{1/[1-\tilde{a}]}.$$

$$E^1 = ([1-\tilde{a}] U^1 / K)^{1/[1-\tilde{a}]}$$

$$= E^0 [U^1/U^0]^{1/[1-\tilde{a}]}$$

$$E^1-E^0 = \{ [1-\tilde{a}] U^1 / K)^{1/[1-\tilde{a}]} - ([1-\tilde{a}] U^0 / K)^{1/[1-\tilde{a}]} \}$$

$$= E^0 ([U^1/U^0]^{1/[1-\tilde{a}]} - 1).$$

The proportional equivalent variation is given by that multiple of original wealth that gives the new utility:

$$U^0 = K (W^0)^{1-\tilde{a}}/[1-\tilde{a}]$$

$$U^1 = K (kW^0)^{1-\tilde{a}}/[1-\tilde{a}]$$

$$= k^{1-\tilde{a}} U^0.$$

$$k = (U^1/U^0) ^{1/[1-\tilde{a}]}.$$

Note that the equivalent wealth variation in terms of k is

$$E^1-E^0 = E^0 (k - 1).$$

When applied to the problem without borrowing constraints, the equivalent variation is equal to the net transfer,

$$E^1 - E^0 = T,$$

and the proportional variation is equal to the ratio of the net transfer to lifetime earnings,

$$(E^1 - E^0)/E^0 = T/Y.$$

Proportional variations, $\tilde{a}=1$ utility, and mortality differences

In comparing equivalent variations between two persons with different lifespans, the shorter lifespan will have a smaller E^0, E^1, and E^1-E^0, other things being equal. It will not necessarily have a smaller proportional variation. In fact, since

$$U = K\, W^{1-\tilde{a}}/[1-\tilde{a}],$$

$$U^1/U^0 = (\, W^1/W^0\,)^{1-\tilde{a}},$$

only that part of the mortality difference that affects W will be present in the proportion.

$$U = \{\, \Sigma_x\, RL_x\, (\, RL_x/DM_x\,)^{-1/\tilde{a}}\, \}^{\tilde{a}}\;\; W^{1-\tilde{a}} \,/\, [1-\tilde{a}]\, .$$

As $\tilde{a}$ approaches 1, the part in pointy brackets simplifies. The other part behaves like log W:

$$U(\tilde{a}=1) = \log (W)\ \{\, \Sigma_x\, DM_x\, \}.$$

The annualized utility, dividing lifetime utility by $\Sigma_x\, DM_x$, therefore approaches

$$U(\tilde{a}=1) \,/\, \{\, \Sigma_x\, DM_x\, \} = \log W.$$

The difference in annualized utility is

$$U^1 - U^0 = \log (W^1/W^0).$$

Borrowing constraints

Capital each period can be calculated as

$$RL_t\, k_t = \Sigma_{x<t}\, RL_x\, (y_x\, [+\, t_x\,] - c_x\,).$$

The capital from the optimal path is checked for negative values. If any are found, consumption is iteratively re-optimized by segment, setting where necessary k_t to 0 and c_t to $RL_{t-1}/RL_t\, k_{t-1} + y_t$.

Quasi-equivalent variations

For the borrowing constraint simulations, an equivalent variation is not calculated. Instead, the equivalent variation formula from the unconstrained problem is used as a utility index. I.e., U^0 and U^1 are calculated under borrowing constraints, and then E^1=E(U^1), E^0 = E(U^0), and E^1-E^0 are calculated using the unconstrained formula applicable under each individual's mortality. In general E^0 under this procedure will not equal Y, but I have assumed, without checking, that E^1-E^0 gives a satisfactory approximation to the variation that would be calculated using constrained expenditures.

Bequest valuation

As described in the main text, the unconstrained path is first calculated with

$$u'(c_t) = â\ R_t/D_t\ .$$

If the terminal capital

$$k_T = 1/R_T\ \Sigma_x\ R_x\ (y_x\ [+\ t_x\ ]\ -\ c_x\ )$$

is negative, then the utility trajectory

$$u'(c_t) = R_t/DM_t\ \ (ì_T + â\ M_t)$$

is solved for that $ì_T$ that makes terminal capital $k_T = 0$. The first derivative

$$dk_T/dì_T = 1/ã\ \Sigma_x\ R_x^2/DM_x\ \{\ R_x/DM_x\ (\ ì_T + â\ M_t)\ \}^{-1/ã-1}$$

is used iteratively to help solve for $ì_T$.

The linear bequest parameter â is given by

$$â = ê_1\ ê_g\ ë*.$$

In these simulations, $ê_g$ is set to 1 and $ê_1$ to either 0.2 or 1.0. The factor ë* is set equal to

$$ë* = ë^0 + á\ (\ ë^0_{avg} - ë^0)$$

where $ë^0$ is the $W_R{}^0$ regular saving multiplier for the person and $ë^0_{avg}$ is the sample average of all such multipliers.

Figure 1A: Log Mortality

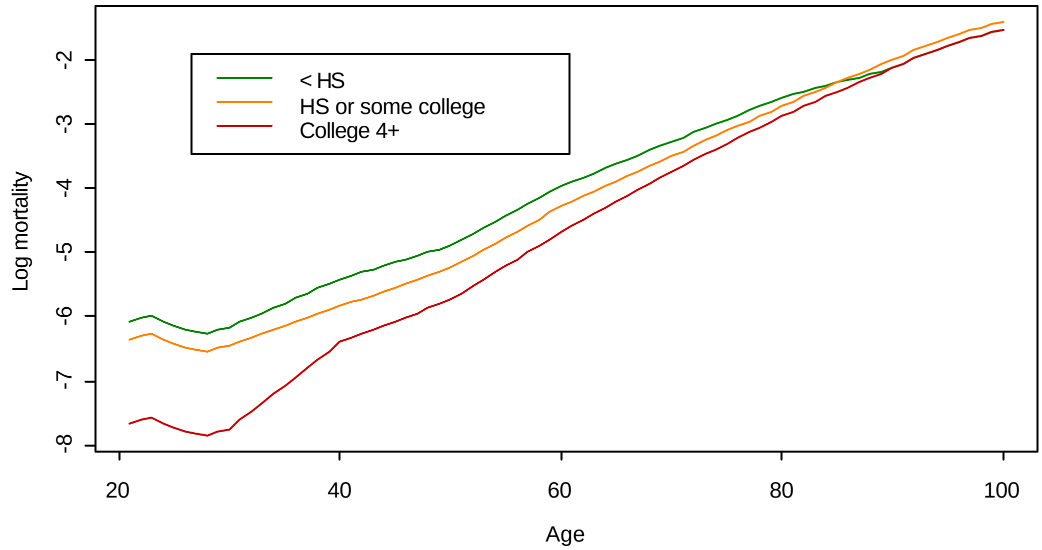

Figure 1B: Survival from age 21

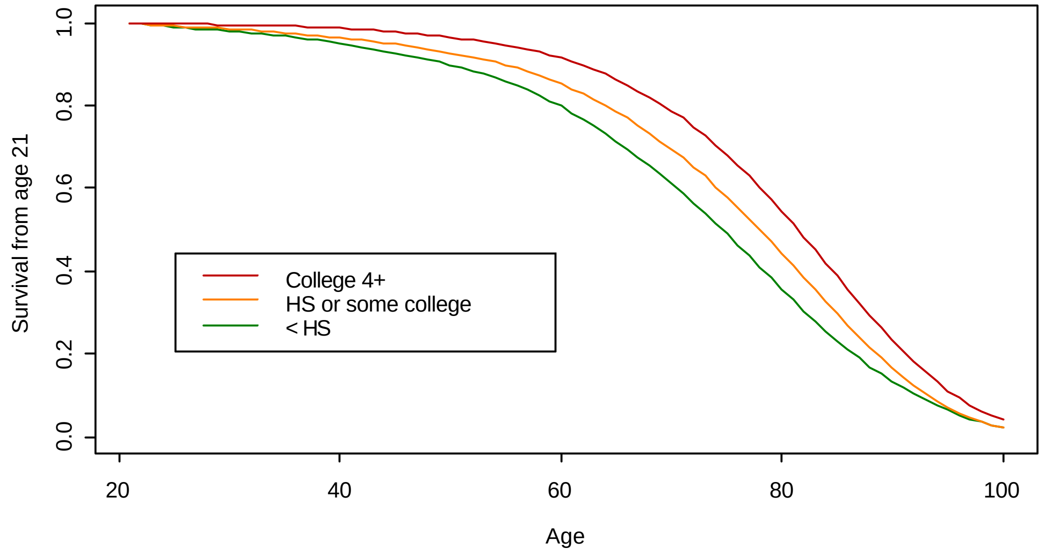

Source: Waldron estimates on 1973 Exact Match, Deaths through 1997

Figure 2: Lifetime values using Differential Mortalities

2A: Lifetime Tax, Benefit, and Net Transfer

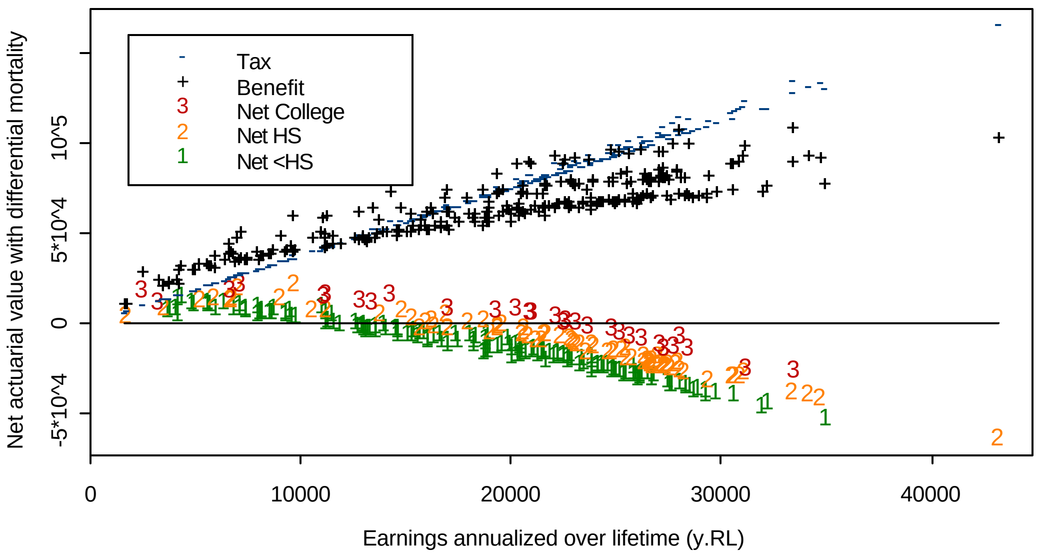

2B: Lifetime Net Transfer

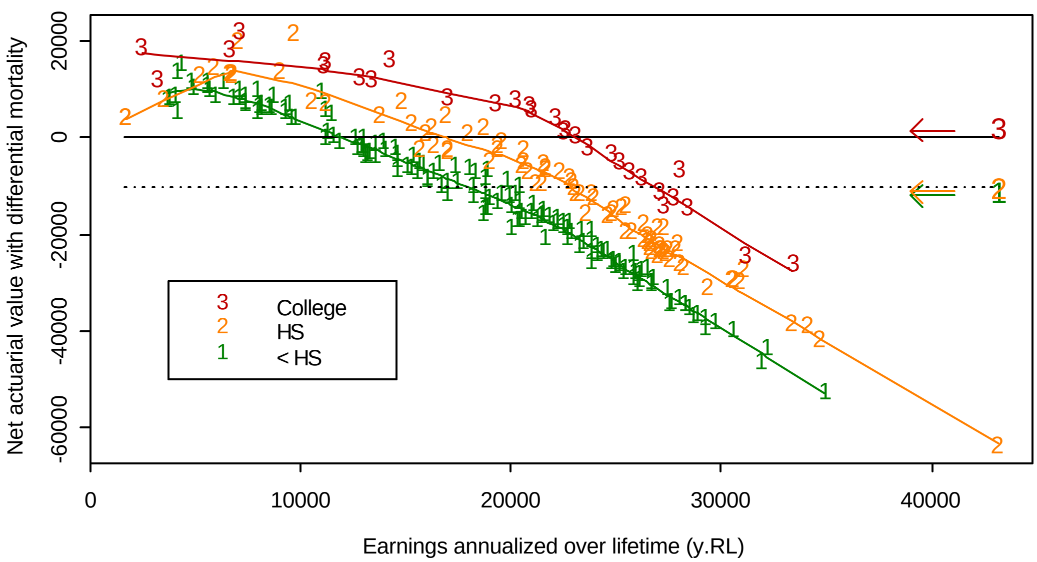

Figure 3: Net Lifetime Transfer

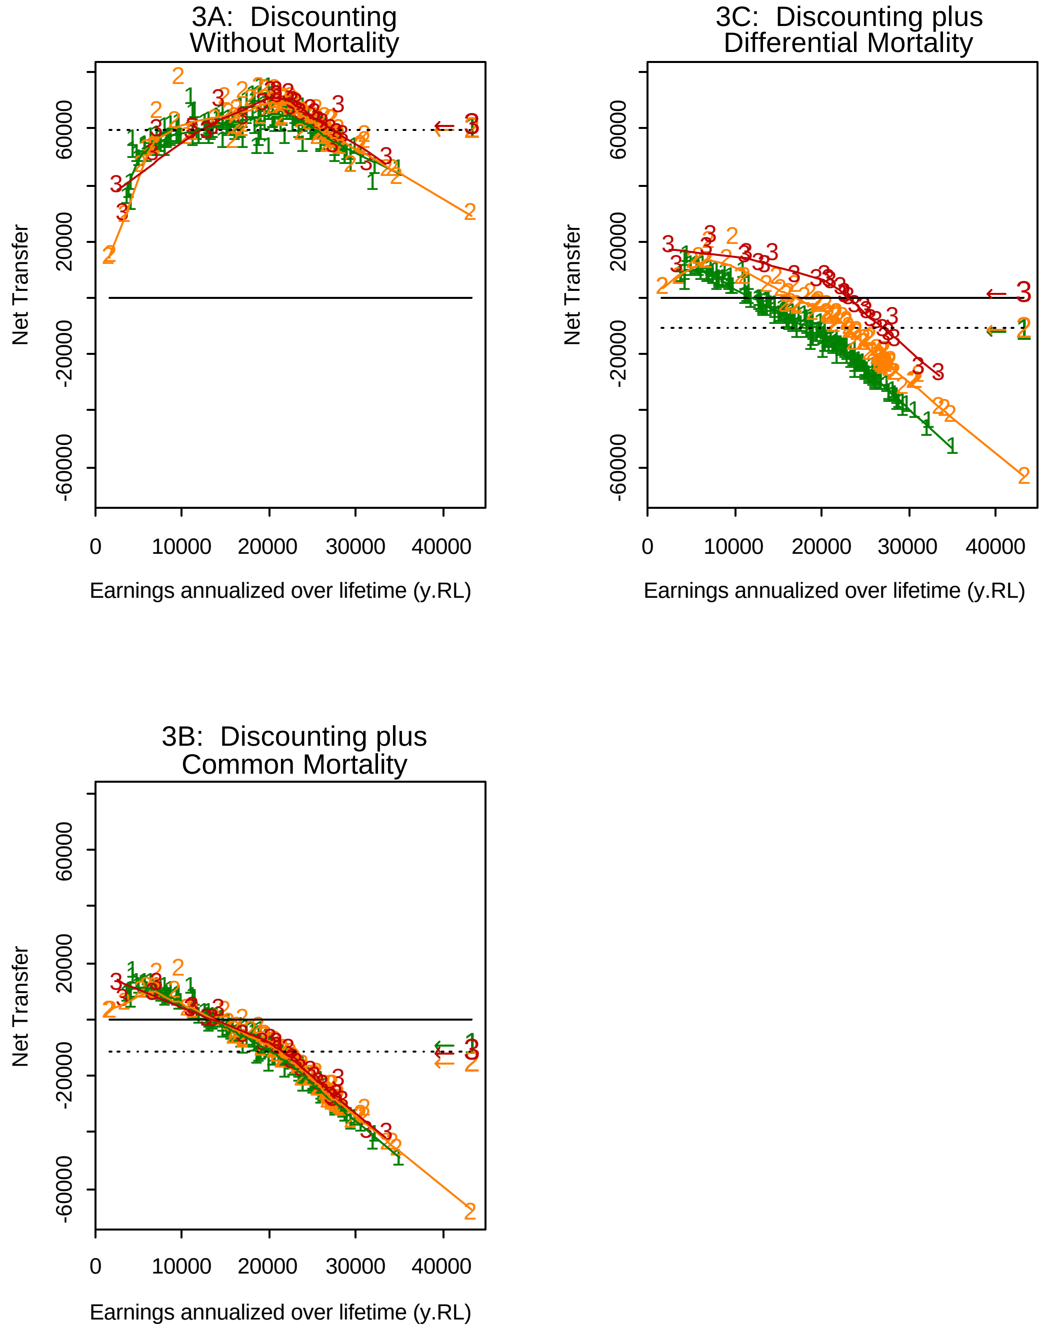

Figure 4: Net Transfers as Proportion of Lifetime Earnings

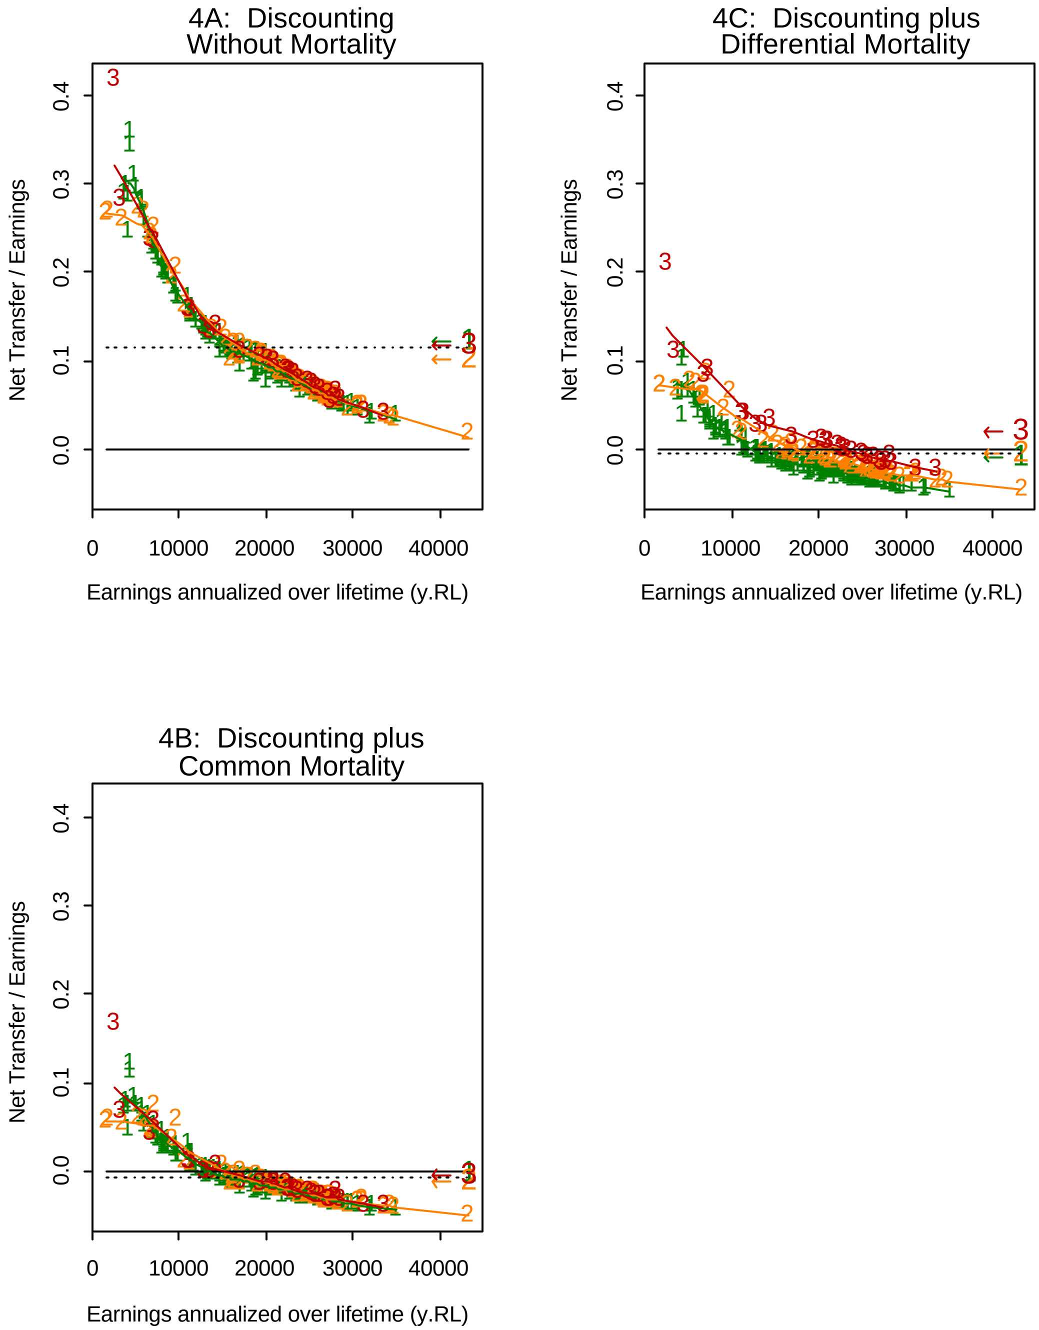

Figure 5: Equivalent variations, risk-aversion coefficient near 1

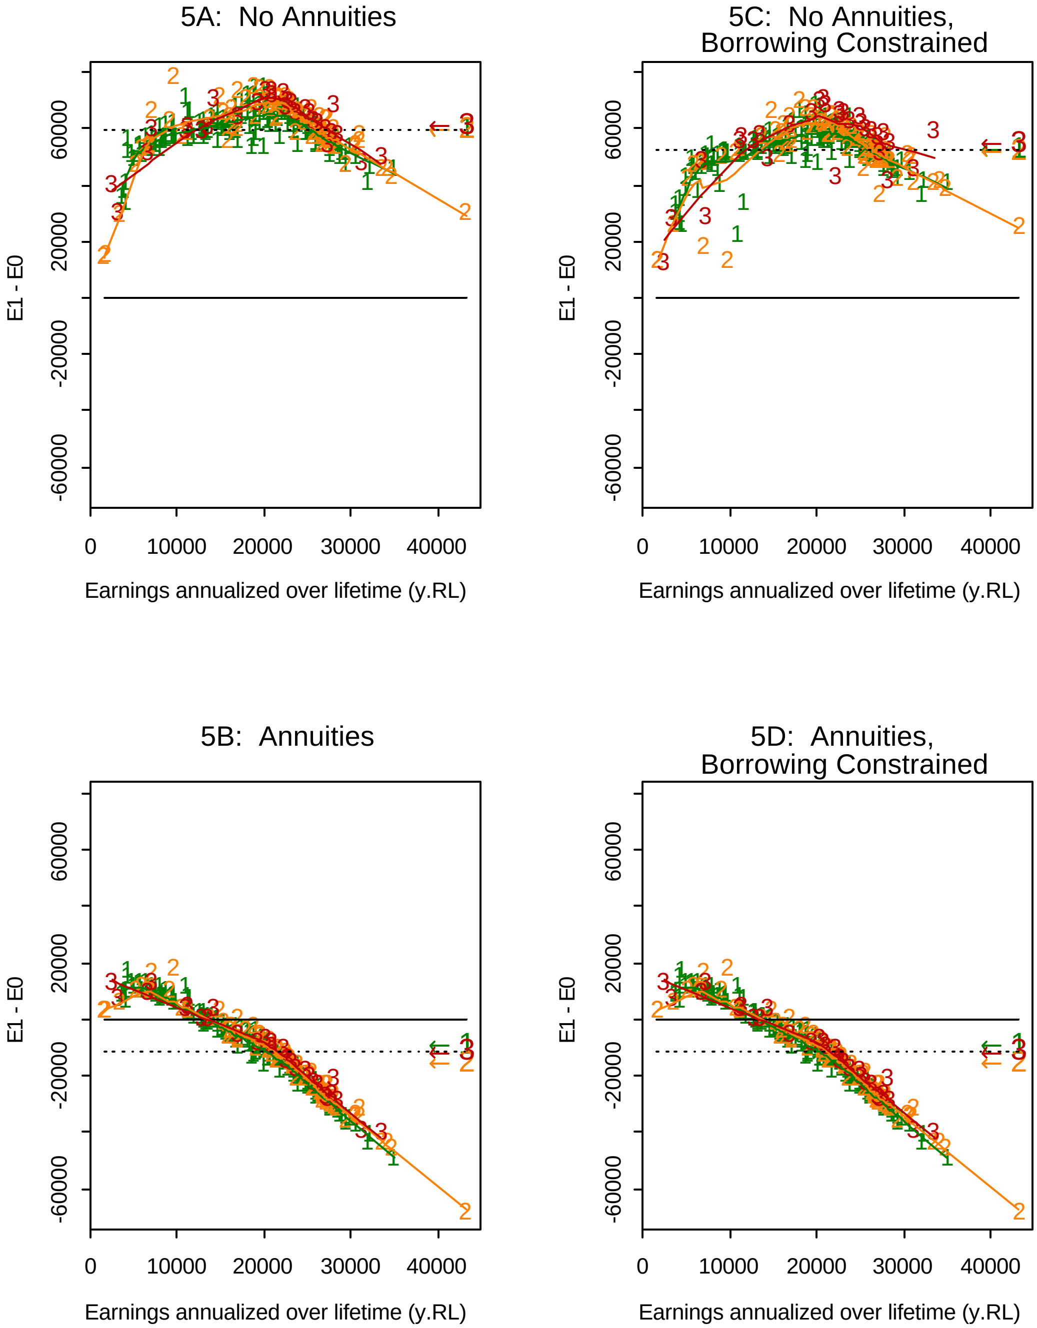

Figure 6: Equivalent proportional variations, risk-aversion coefficient near 1

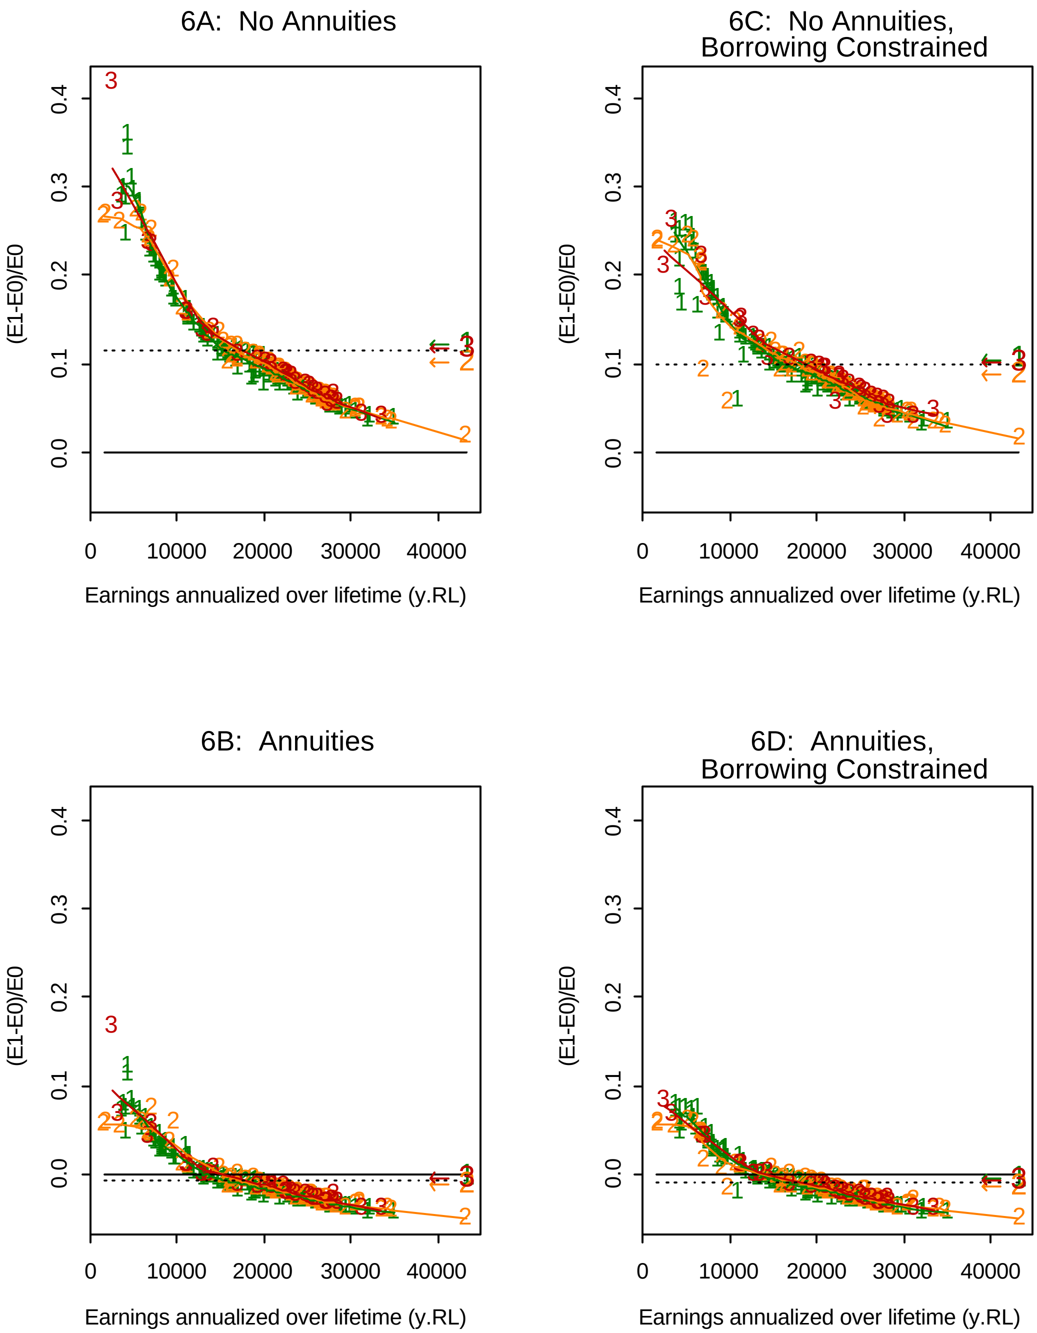

Figure 7: Lifetime utility change, risk-aversion coefficient near 1

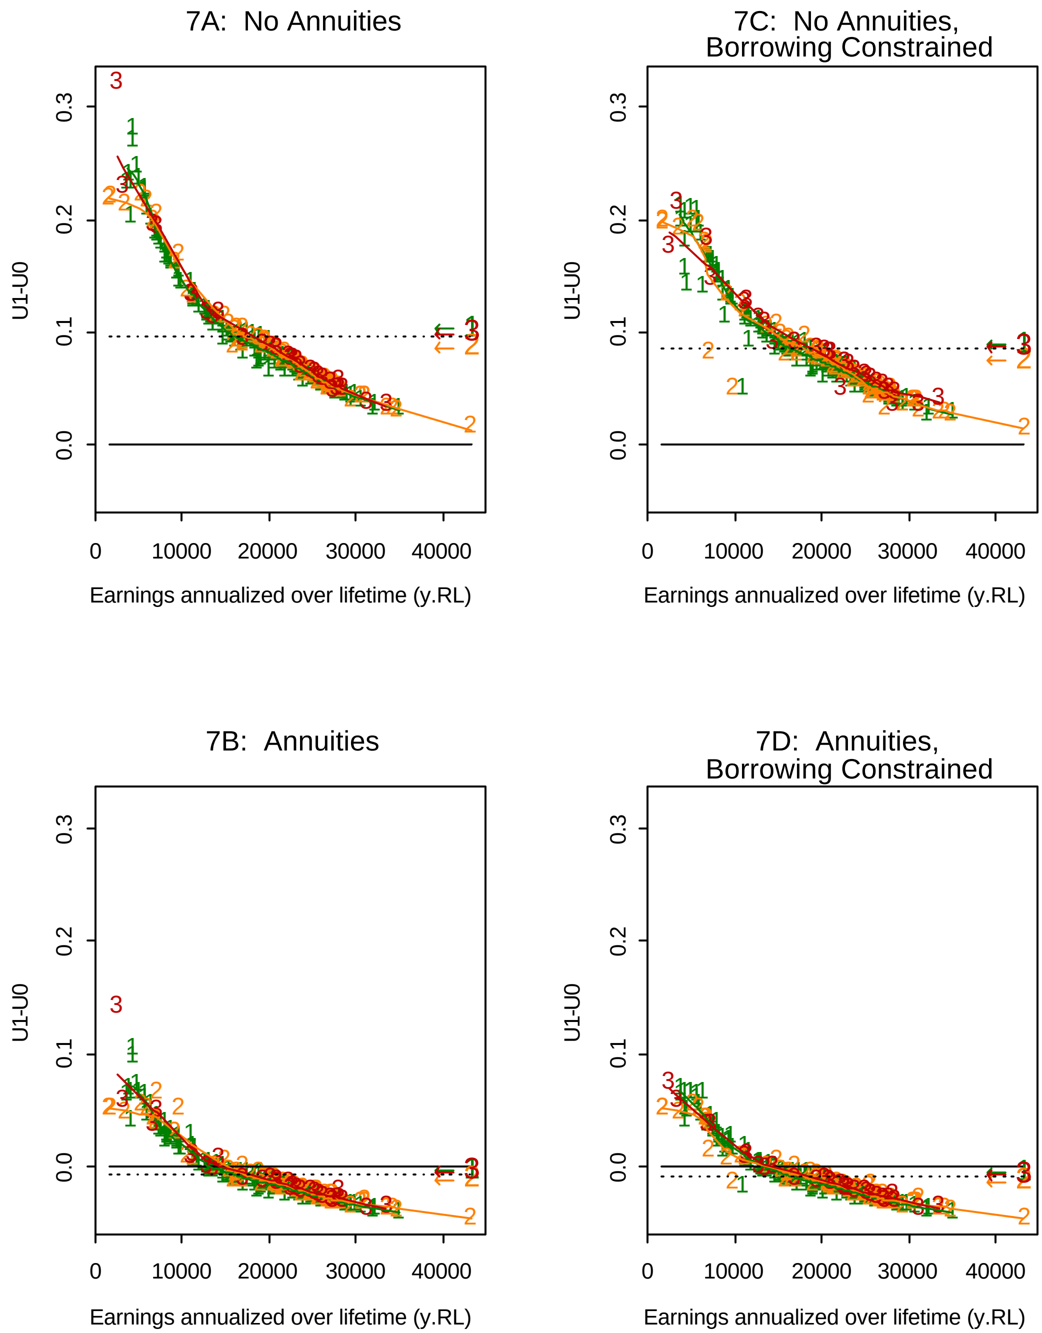

Figure 8: Effect on Utility Change (Unconstrained Annuitization)
Of Changes in Gamma

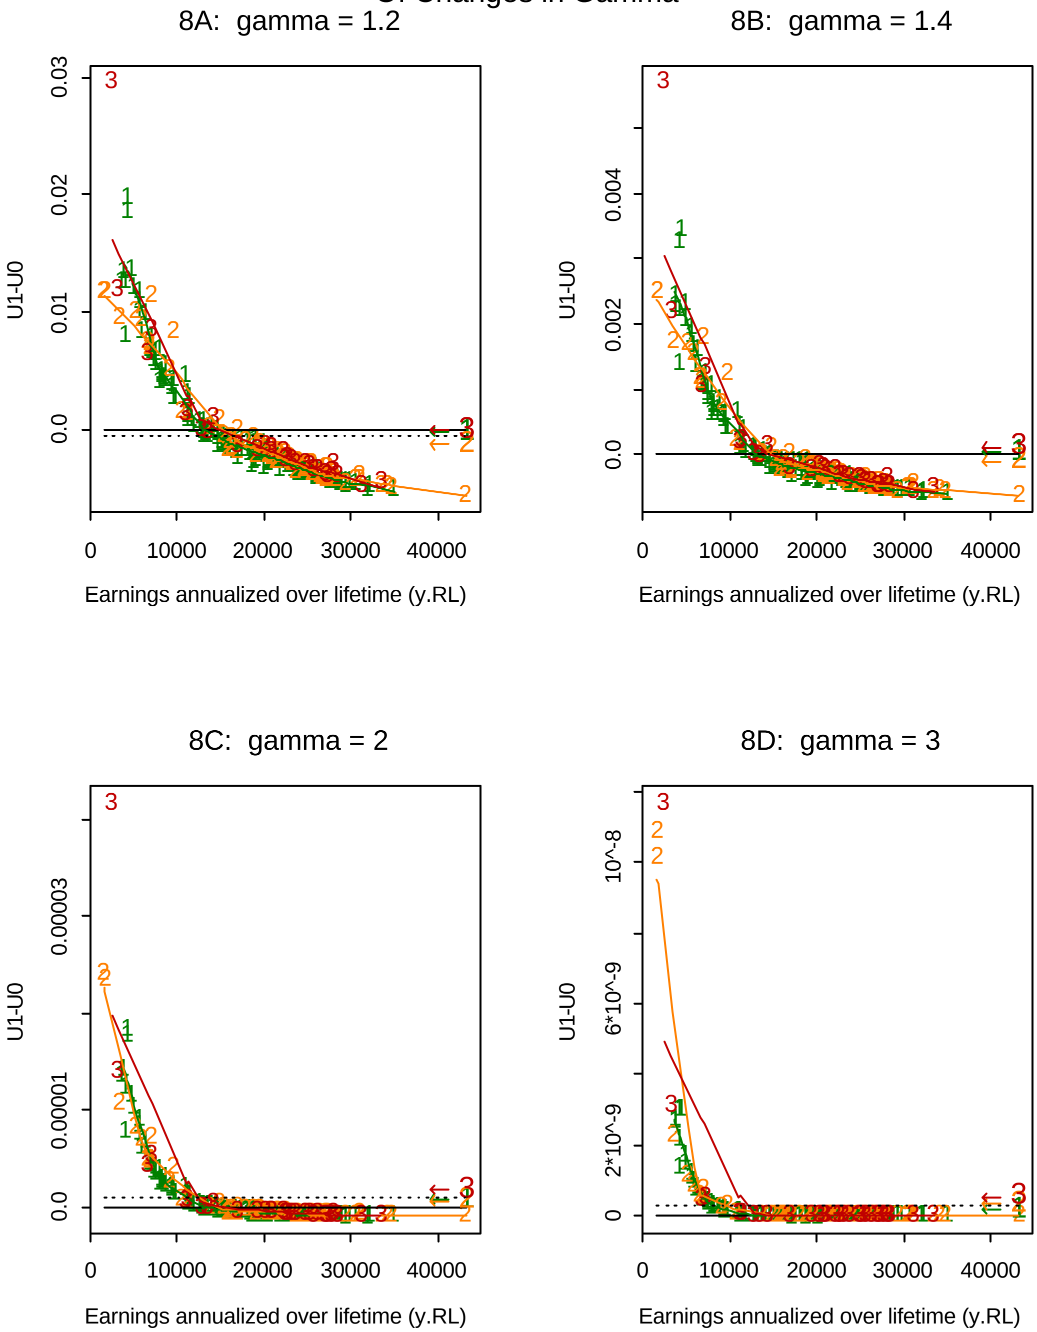

Figure 9: Lifetime utility change with bequest motive, risk-aversion coefficient near 1

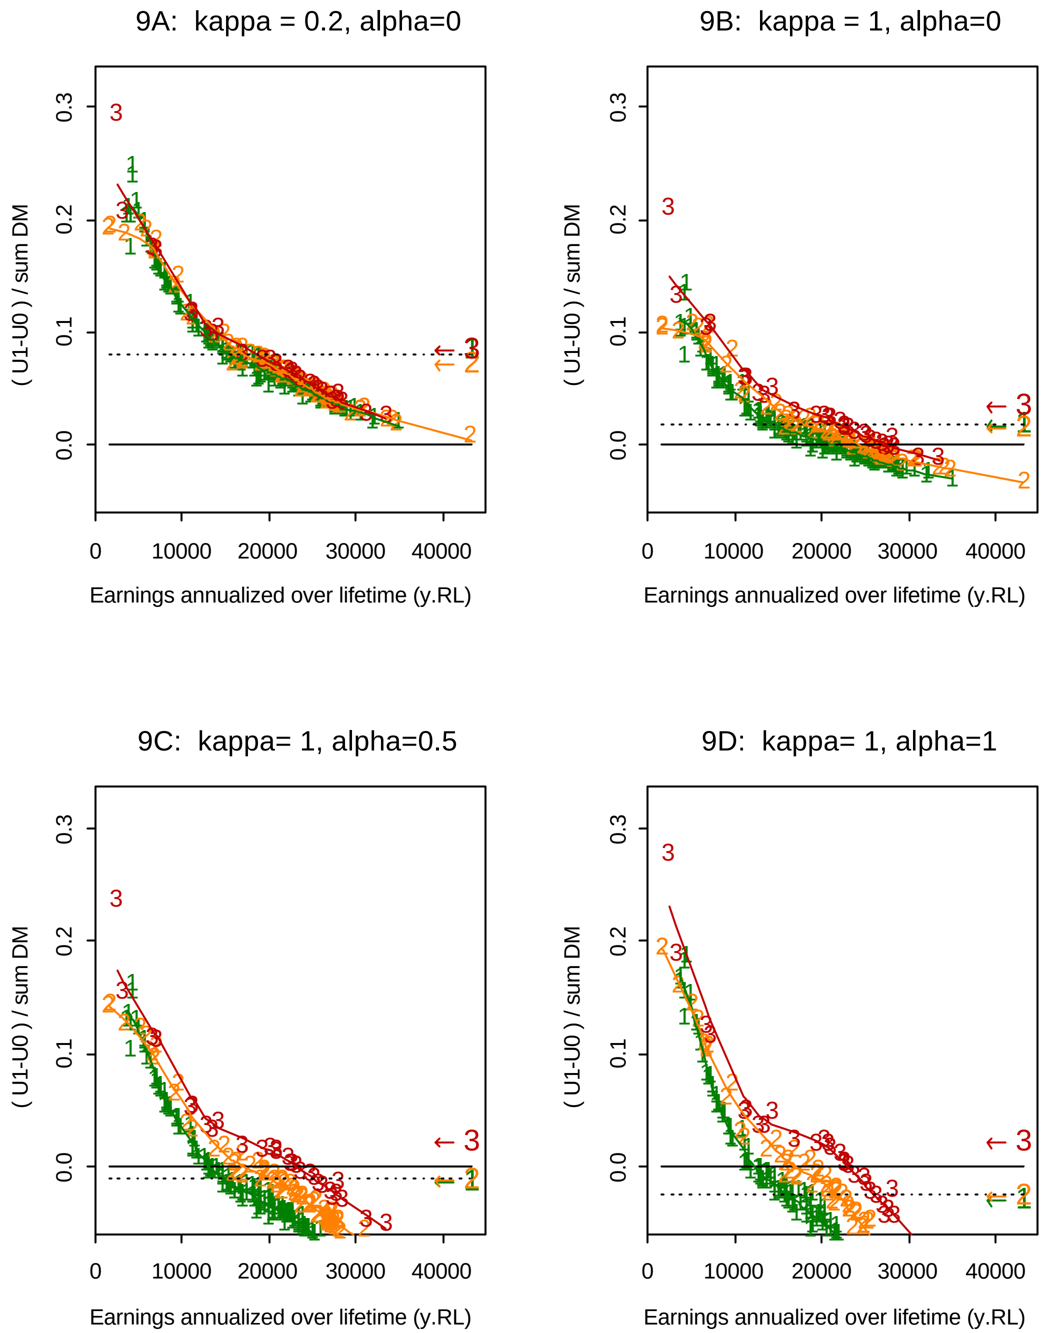

Kappa is bequest-strength parameter (0 = no bequest motive, 1 = perfect altruism).
Alpha: 0 = bequest based on own wealth, 1= bequest based on population average wealth.

Figure 10: Lifetime Uc change with bequest motive, risk-aversion coefficient near 1

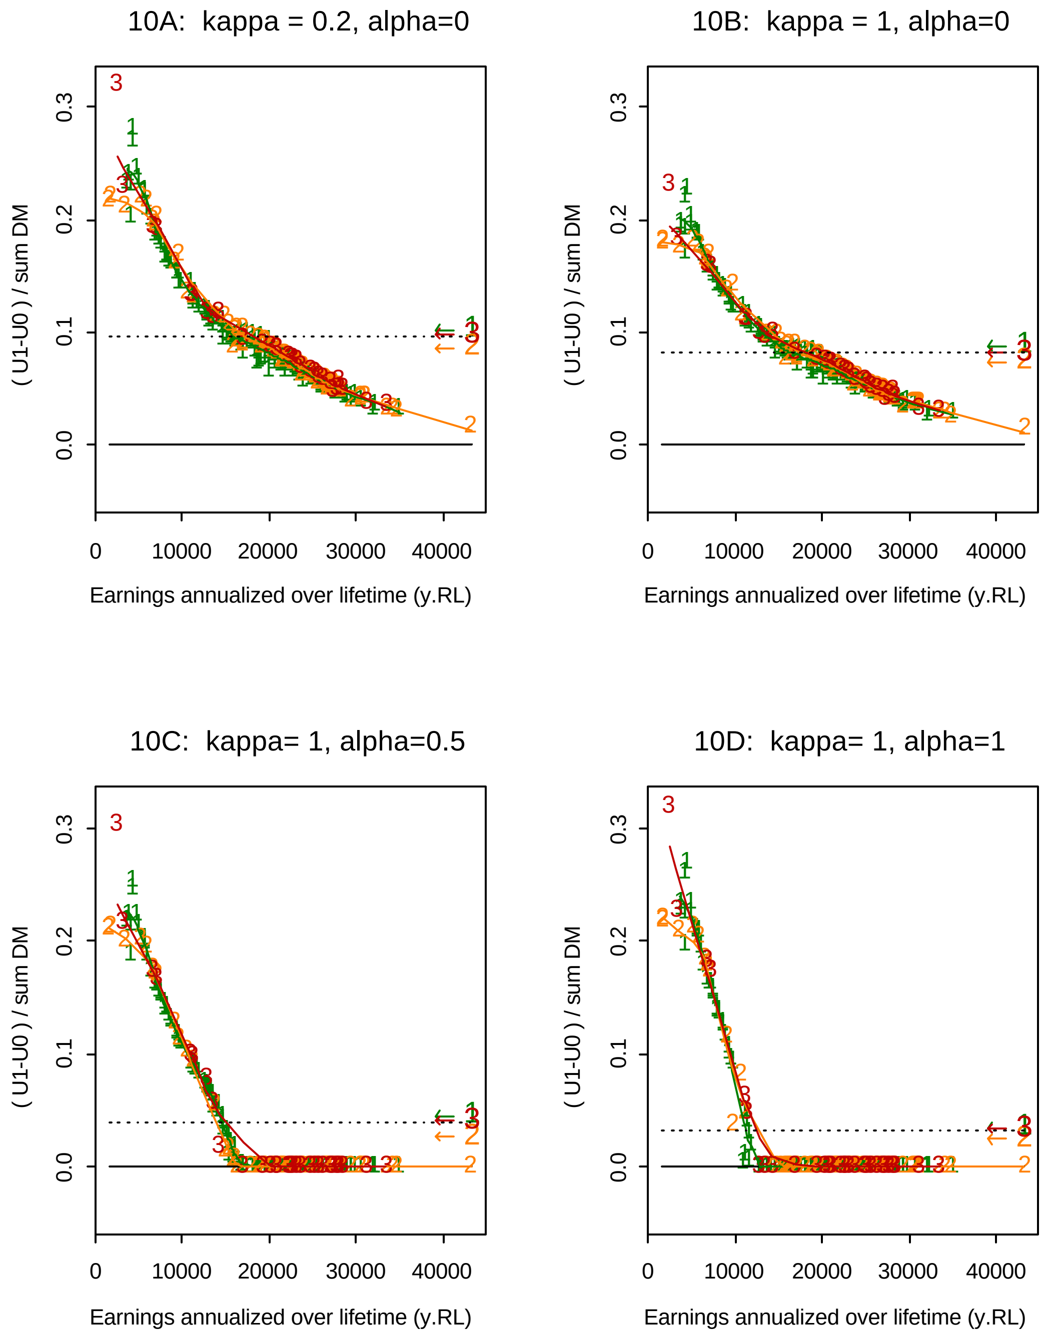

Kappa is bequest-strength parameter (0 = no bequest motive, 1 = perfect altruism).
Alpha: 0 = bequest based on own wealth, 1= bequest based on population average wealth.