HOW DOES DIFFERENTIAL MORTALITY AFFECT SOCIAL SECURITY FINANCES AND PROGRESSIVITY?[1]

Amy Rehder Harris (e-mail: amy.harris@cbo.gov)
John Sabelhaus (e-mail: john.sabelhaus@cbo.gov)

Congressional Budget Office
Washington, DC

March, 2005

Abstract

The Congressional Budget Office Long-Term (CBOLT) model uses dynamic micro-simulation for a representative sample of the population to analyze the aggregate and distributional effects of Social Security policy. In the model, overall mortality rates by age and sex are calibrated to match Social Security Trustees projections, and differential mortality (the difference in death rates across socioeconomic groups) is introduced using a combination of disability-specific mortality and a technique for the non-disabled developed by Lillard and Panis (1999). In this paper the question of how differential mortality affects Social Security finances and progressivity is approached through sensitivity analysis using CBOLT. The model is solved using a range of assumptions about differential mortality, and the impact on various outcomes is assessed. In contrast to inferences in several recent studies, the results here suggest that differential mortality does not play a significant role in determining progressivity or system finances. It is true that socioeconomic differentials in death rates work counter to Social Security's statutory redistribution and make the system costs higher, but the effects are probably only of second-order significance.

---

[1]The views expressed here are those of the authors and should not be interpreted as those of the Congressional Budget Office. We would like to thank Josh O'Harra and Kevin Perese, who contributed to development of the Congressional Budget Office Long-Term (CBOLT) modules used to produce the results for this paper.

1. Introduction

It is widely acknowledged that the U.S. Social Security system is expected to become insolvent within a few decades (Social Security Administration (2004), Congressional Budget Office (2004-A)) and also clear that any steps taken to address the expected financial imbalance will involve a decrease in net benefits (a reduction in benefits or increase in taxes) for some group at some point in time. This tradeoff between restoring overall financial solvency and avoiding negative distributional consequences is at the heart of the debate over how and when to reform Social Security. One of the modeling assumptions that affects estimates of both system finances and progressivity is differential mortality–the principle that socioeconomic status and mortality rates are negatively correlated. This paper uses the Congressional Budget Office Long-Term (CBOLT) micro-simulation model to investigate the role that differential mortality plays in determining overall system finances and progressivity.

Differential mortality has been raised as an argument by various authors to support very different visions for reforming Social Security. The current Social Security system is progressive (in a statutory sense) because beneficiaries with low lifetime earnings get a higher replacement rate (ratio of benefits to average lifetime earnings) than their higher-earning counterparts. However, most researchers agree that the statutory redistribution is offset at least in part because beneficiaries with low lifetime earnings die younger and get benefits for a shorter period. Some proponents of moving the system towards individual accounts have argued that–because the current system is not effectively progressive anyway–having net benefits tied more closely to individuals' earnings would not affect (or could even increase) overall progressivity.[2] At the other end of the political spectrum, some proponents of raising net taxes on higher lifetime earners argue that the goal of reform should be to establish (or re-establish) progressivity in a system where it is currently lacking because differential mortality eliminates the desired redistribution embodied in the law.[3] As a corollary, if the long-run financial imbalance in Social Security is attributable to higher lifetime earners collecting benefits longer, then it is argued higher earners should be the ones to fund those benefits.

The idea that socioeconomic status and mortality are negatively correlated is well established in a number of studies, and there is some evidence that the relationship is changing (at least by education) over time.[4] This paper does not weigh in on that debate, focusing instead on using sensitivity analysis as an approach to understanding the role played by differential mortality in Social Security. The CBOLT micro-simulation model introduces differential mortality through two distinct channels, and those channels can be altered to test for the significance of differences in death rates. First, the model incorporates a direct relationship between earnings and Disability Insurance (DI) claiming and then between DI status and mortality, generating an indirect (but very powerful) link between earnings and mortality. Second, for the non-disabled, CBOLT adopts the Lillard and Panis (1999) model of differential mortality which has explicit socioeconomic covariates including earnings, marital status, and

---

[2]See, for example, Feldstein (2005).

[3]See, for example, Diamond and Orszag (2004).

[4]See Preston and Elo (1995), Pappas, Hadden, and Fisher (1993), Lauderdale (2001), and Elo and Smith (2003).

education.[5] This second channel in particular allows some flexibility (by varying estimated coefficients) for specifying the effect that differential mortality has on progressivity and system finances, and thus it is possible to run counter-factual experiments to understand the role played by differential mortality.

The effect of differential mortality on Social Security system progressivity has been analyzed in several previous papers, but the results here are an important step forward for several reasons. First, most previous studies focus only on the non-disabled population, and generally only hypothetical "example" earners or one cohort of individuals who reach (or are projected to reach) retirement benefit eligibility age. CBOLT operates on a representative sample of the population for the period between 1984 and (in a typical simulation) 2110, and thus allows much broader cross-sectional and longitudinal analysis. The second confounding factor in previous studies of progressivity is that differential mortality is often analyzed along with other variables that affect estimates of redistribution–most notably, the characteristics that are used to classify individuals into the groups across which redistribution is being measured.[6] Indeed, some researchers have concluded that Social Security in its current form is not progressive at all, once individuals are properly classified.[7] That conclusion is often interpreted to mean "differential mortality eliminates the statutory redistribution in Social Security." Indeed, it is shown here that

[5]The Lillard and Panis (1999) differential mortality formulation was originally developed for the Social Security Administration's Modeling Income in the Near Term (MINT) micro-simulation. The MINT model has been used extensively is previous studies of Social Security redistribution–see, for example, Cohen, Steuerle, and Carasso (2001) and Smith, Toder, and Iams (2003/2004).

[6]Fullerton and Mast (2005) organize their discussion of Social Security progressivity around this classifier principle.

[7]Coronado, Fullerton, and Glass (2002), Liebman (2002), Gustman and Steinmeier (2001), Armour and Pitts (2004).

differential mortality probably plays a fairly modest role in that conclusion once the differences in classifier variables are accounted for.

In addition to studying progressivity, the sensitivity analysis approach used in this paper is also useful for investigating the role that differential mortality plays in Social Security system finances. The experiments presented here are designed to address how much of the long-run financial imbalance in Social Security is attributable to the relative differences in life expectancy suggested by various models of differential mortality. As with the conclusions about progressivity, it is hard to argue that (holding overall mortality by age and sex fixed) differential mortality has a first-order impact on system finances.

These conclusions about the effect of differential mortality on Social Security finances and progressivity are both model and measure dependent, but there is good reason to believe they would generalize to other models and alternative measures for outcomes or classifiers when subjected to the same sort of test. The reason to expect that the conclusions about differential mortality will generalize is quite simple: differences in death rates across age groups dominate the effect of differential death rates within age groups. That is, even when the extent of differentials is increased arbitrarily to generate very skewed death rates by income at every age there is only a modest effect on conclusions about overall progressivity or finances, because differences in death rates across age groups are so much larger. The simulations here hold overall mortality by age and sex fixed, and that effect is shown to dominate the results.

2. Differential Mortality in the CBOLT Micro Model

Micro-simulation is a powerful tool for analyzing public policy issues when the interaction of complex program rules and population heterogeneity is likely to be of first-order significance–Social Security is a prime example of a program about which micro-simulation is well suited to provide important insights. The Congressional Budget Office Long-Term (CBOLT) model integrates the micro-simulation approach with a macro-economic/unified budget framework, making it particularly useful for studying questions that involve both aggregate and distributional analysis. This section provides an overview of CBOLT, and describes the methods by which differential mortality is introduced into the analysis.

Overview of the CBOLT Model

The basic design principle behind the CBOLT micro-simulation is to first generate realistic demographic, economic, and policy outcomes for a large representative sample, then apply the complex Social Security program rules and solve for micro-distributional and aggregate budgetary outcomes. The modeling process begins with a large-scale Social Security Administration (SSA) longitudinal data file from which a sample is drawn to become the basic input to the micro simulation. The sampled data file is then used as a starting point for applying CBOLT's micro-processes (methods for projecting individual demographic and economic outcomes over time) which are estimated using various data sources and calibrated over history. The micro simulation is also integrated with a comprehensive macro/budgetary model, so the projections provide insights about both Social Security system finances and estimates of progressivity.[8]

The starting point in any micro-simulation model is the base data file. CBOLT uses the Social Security Administration (SSA) Continuous Work History Sample (CWHS), which is a one-in-one-hundred sample covering every Social Security number ever issued. For each observation in the data file, the CWHS reports a comprehensive earnings and worker benefit history along with basic demographics. The primary advantages of the CWHS are that the sample is very large (indeed, CBOLT uses only one-tenth of the CWHS, so the model itself is a one-in-one-thousand sample) the data is from high-quality administrative records (as opposed to limited and sometimes biased self-reported data) and the data set is updated annually (so CBOLT can be re-based every year). The primary disadvantage of the CWHS is the lack of comprehensive demographics and other information that is available from public surveys, but that shortcoming is resolved in the calibration process discussed below.

Given a micro base file, the second step in any micro-simulation is specifying transition equations–the processes by which demographic and economic outcomes for individuals in the micro sample evolve over time. CBOLT operates on the basic processes (birth, education, labor supply, earnings, first marriage, divorce, remarriage, mate matching, benefit claiming, benefit awards, and ultimately death) needed to calculate Social Security taxes and benefits and to integrate the micro outcomes with the macro growth model and unified budget framework.[9]

---

[8]The macro growth model framework is not discussed here, but is closely related to the approach in Bosworth and Burtless (2002, 2004) with two important differences: aggregate labor input is summed from the micro model in CBOLT, and aggregate private saving adjusts to target a stable long-run capital output ratio. That saving assumption effectively neutralizes the impact of assumptions about other components of the Federal budget and creates a stable baseline similar to the one used by the Social Security Trustees when they analyze system finances.

[9]The micro transition processes are described in detail in a series of technical papers available on the CBO web site. For a more detailed overview of the micro-modules, see O'Harra, Sabelhaus, and Simpson (2004); for a discussion of the marital transition processes see O'Harra and Sabelhaus (2002); for a discussion of the labor force

Those processes are estimated using a number of different data sources that each provide unique information on a particular demographic or economic process. The micro-transition processes are primarily designed to capture correlations between individual characteristics (as opposed to reflecting optimizing economic decisions) and thus the results are useful for distilling the first-order impact of choosing micro-simulation for studying Social Security.[10]

The micro transition processes developed for projecting future individual outcomes are also used to assign information that is not available in the base (CWHS) data file. CBOLT uses a "historical simulation" approach to assign the demographic characteristics that are not present on the micro base file. Although this type of (imputation-based) assignment is used in all major micro-simulation projects, it is quite extensive in CBOLT and deserves special mention. The basic idea is to assign missing characteristics in history using the same methods for projecting forward, then test and calibrate the processes using external data sources available in history.

The initial CBOLT micro sample is actually drawn to be representative of the U.S. population for the period 1984 through 2002.[11] The historical simulation process begins in 1984, when each individual gets initial unobserved characteristics assigned based on their observed characteristics (a standard imputation). But then, each micro process is applied and then

participation and earnings projections see Harris and Sabelhaus (2003); and for a discussion of the mate-matching algorithms see Perese (2002).

[10]Even the limited-behavior version of the micro-simulation based-approach leads to important insights about Social Security that do not come through in other analyses. First, all else equal, the micro-simulation generates projected benefit awards for male OAI workers below those based on standard actuarial techniques, because CBOLT properly captures observed shifts in the historical relative earnings profiles (CBO, 2004-A). Second, direct analysis of the micro-level outcomes suggests there are serious problems with using hypothetical "example" workers to analyze the impact of proposed reforms (CBO, 2004-B, C, D).

[11]Social Security coverage rates were much lower for working age cohorts before then, and also the availability of earnings data above the taxable maximum did not occur until the 1980s.

calibrated so it generates the actual observed distributions (say population by marital status) that are known from some external data sources. The model is also carefully tested for its ability to reproduce empirical covariances (say between husband and wife ages). Perhaps the most important test of the model is that it matches Social Security system outcomes (numbers of beneficiaries and average benefits by type of benefit award) in the historical period.

How Differential Mortality Enters in CBOLT

Differential mortality is introduced into CBOLT through two channels. The first channel is for the disabled (actually, the Disability Insurance (DI) beneficiary) population, for whom differential mortality is implicit because the probability of becoming a DI beneficiary is a function of earnings, and the probability of a transition off DI through death is much higher than death rates for the non-disabled population. The second channel for differential mortality applies to the non-disabled, and uses estimates of the correlation between death rates and various socioeconomic characteristics from Lillard and Panis (1999).

Before discussing the two distinct mortality processes in detail, it is useful to step back and understand how differential mortality fits into the general flow of a CBOLT simulation (see Figure 1). The simple schematic in Figure 1 for differential mortality over the course of one simulation year suggests some important connections that should be kept in mind. Working backwards, the actual period t mortality covariates for the disabled population are age, sex, and length of time on DI. For the non-disabled, the covariates again include age and sex, but also the common variables between CBOLT and the Lillard and Panis (1999) equation, which are education, marital status, and earnings. The schematic makes it clear that the various paths for differential mortality causation are actually more complicated than the final covariates would suggest. For example, in CBOLT, earnings affect the probability of marital transitions (O'Hara and Sabelhaus, 2002) which in turn affects labor supply (Harris and Sabelhaus, 2002) which in turn affects earnings in the next period. Thus, even if marital status was the only covariate (other than age and sex) in the mortality equation, there would still be differential mortality in the model (with respect to earnings) through the other linkages.

The link between earnings and mortality for the DI population in CBOLT is very straight-forward. The probability of entering DI in any given year is a function of age, sex, and earnings. The effect of earnings is estimated using the CBOLT input data file (the CWHS, described above) for the period 1984 to 2002 (see Figure 2). The bar chart shows the fraction of new DI entrants coming from each earnings decile–workers in the bottom decile have something like four times the probability of going on DI relative to workers in the top decile. As with other aspects of the model, these probabilities are evaluated in the "historical simulation" process with an eye towards whether or not the distribution of benefit awards is consistent with actual patterns.

Given that an individual in CBOLT is assigned DI beneficiary status, they face a much higher death rate than if they had remained non-disabled (see Figure 3). The probability of death for DI beneficiaries in CBOLT is based directly on data provided by the Office of the Chief Actuary (OCACT) at the Social Security Administration. The death rates for new DI beneficiaries are quite high for the first two years a person is receiving benefits, but even the rates for people who have been on DI for ten years are very high–orders of magnitude larger than the average death rates in these age/sex groups. The combination of earnings-based entry onto

DI and much higher than average death rates together generate an indirect but very powerful differential mortality in the outcomes.

The differential mortality process for the non-disabled population is based directly on the work of Lillard and Panis (1999).[12] In order to estimate a differential mortality equation for the SSA Modeling Income in the Near Term (MINT) project, Lillard and Panis (1999) used the Panel Survey of Income Dynamics (PSID) data set that provided them with both cross-sectional and longitudinal variation. The authors used a number of covariates that are also included in the CBOLT model (education, marital status dummies, and earnings) but also a few that are not (race, in particular). In addition, they went to great lengths to adjust the estimated coefficients on the basic variables (age, sex, and calendar time) so that the PSID-based estimates generated overall death rates that matched Vital Statistics outcomes.

The CBOLT differential mortality module uses a subset of the coefficients from the Lillard and Panis (1999) model because (1) only that subset of the socioeconomic independent variables overlap, and (2) CBOLT is calibrated to match OCACT projected populations, and thus the basic variables in the equation (age, sex, time) are not needed.[13] Calibration to OCACT implies that the overall mortality rates by age, sex, and year in any simulation are taken as given in a CBOLT simulation–the only question is how those deaths should be distributed within the micro sample in the group. Thus, the approach used for introducing the differential effects is to start with the coefficients from Lillard and Panis (1999), then, during the simulation, adjust the

[12]For an alternative approach to assigning differential mortality using external data sources, see Brown, Liebman, and Pollett (2002)

[13]While developing the CBOLT micro-simulation capability over the last few years CBO has explicitly avoided generating projections based on demographic assumptions that differ from OCACT, focusing rather on the economic differences between the models.

(implied) constant of the log-hazard equation in each age/sex/year cell for the effects of the other independent variables such that the overall predicted death rate for that cell matches OCACT.[14]

The approach taken to introducing differential mortality for the non-disabled captures the effects of education, marital status, and earnings, but there are reasons to be cautious about exactly how to apply the Lillard and Panis (1999) coefficients in CBOLT.[15] There is some (indirect) evidence that CBOLT generates the right amount of differential mortality, at least by earnings, which is the focus of this paper.[16] In any case, the approach here of using sensitivity analysis (which involves arbitrarily raising or lowering the Lillard and Panis (1999) coefficient on earnings) is the best way to approach this situation in which there is a good deal of uncertainty about the size of the differential mortality effect. One of the principal conclusions of this paper is that differential mortality does not seem to affect outcomes as much as other authors have argued–that conclusion is based on varying the impact of earnings across a wide range, thereby generating differences in death rates that vary from negligible to very skewed.

---

[14]The actual process in a micro-simulation is actually a little more complicated because the model operates on a micro sample. In a simulated population sampled at less than 100%, applying expected death rates directly will always generate too little population over long periods, because there is an asymmetry that arises when death is assigned randomly. Basically, you can always generate more deaths when people (randomly) did not die in an earlier year–when they should have–but you cannot bring people back to life who (randomly) died too early. Thus, the actual CBOLT micro algorithm adjusts death rates for cumulative deaths on a cohort basis.

[15]For example, although education and marital status are simple categorical variables, the CBOLT education classes do not match Lillard and Panis (1999) exactly. Also, earnings is a continuous variable that requires some construction. The CBOLT measure is the difference between the individual's and the age/sex mean of "average lifetime earnings" through the year prior to when death is being assigned, which is very similar to the variable Lillard and Panis (1999) constructed with the PSID when attempting to match the comparable variable in the MINT Survey of Income and Program Participation (SIPP) data base.

[16]For example, one can look at how the average lifetime earnings of a cohort changes as they age to gauge whether or not the differentials in death rates make sense, because attrition of below average individuals will increase the average of those remaining. In CBOLT, when differential mortality is assigned using the Lillard and Panis (1999) coefficients, the projected within-cohort average Primary Insurance Amount (PIA)–which is based directly on average earnings–rises systematically between the ages of 62 and 80at roughly the same pace evident in both history and in OCACT projections.

3. Differential Mortality and Social Security System Progressivity

If socioeconomic characteristics and mortality are negatively correlated Social Security will be less progressive in practice than it is in a statutory sense.[17] Several previous studies of Social Security progressivity have been undertaken with differential mortality as one of the considerations in the analysis, rather than the focus, as in this paper. However, those previous studies also suggest that the other important inputs into any estimate of progressivity–the measure of redistribution itself, the population subgroup being studied, and the characteristics used to classify people for measuring net returns–all potentially interact with the effects of differential mortality. Thus, in addition to focusing on the effects of differential mortality, the analysis here presents measures of progressivity using CBOLT which vary all of these inputs as well.

CBOLT Base-Case Estimates of Social Security Progressivity

The starting point for analyzing the impact of differential mortality on progressivity is the published analysis based on the CBOLT model (see Figure 4; from CBO 2004-A). Given the nature of official publications (showing extensive variation for a set of results is less feasible than in an academic publications) the CBO had to make choices from the myriad of possible measures of redistribution, possible population subgroups, and possible classifiers for organizing the sample into the groups across which progressivity is measured. Given those choices, the

---

[17]The statutory progressivity stems from the fact that, although the Social Security tax rate is equal across earnings groups up to the taxable maximum, the formula for determining the actual benefit at retirement (through the Primary Insurance Amount (PIA)) replaces a significantly higher fraction of earnings for low lifetime earners than it does for high lifetime earners. CBO 2004-A shows those replacement rates differ by a factor of two between the lowest and highest lifetime earnings quintiles.

estimates published by the CBO are somewhat striking in that they seem contrary to the common belief (based on recent empirical studies) that Social Security is not very progressive in practice. CBO reported (and it is evident in Figure 4) that Social Security generates a wide disparity in net returns across earnings groups, and that holds for cohorts for whom significant amounts of actual data already exist (the 1940s birth cohort in particular) as well as for cohorts for whom all of their outcomes are projected (1990s and beyond).

The estimates shown in Figure 4 are particularly striking because there is no immediate reason (based on the literature) to suspect that the choices made to generate the estimates are particularly biased towards finding progressivity. The measure of net returns presented is the ratio of the present value of lifetime Social Security benefits to lifetime Social Security taxes, discounted at a fixed real interest rate. If that present value ratio equals 100 percent, then (at the chosen discount rate) the people in that group get out (in present value) exactly what they put in. Thus the estimates in Figure 1 suggest individuals in the lowest earnings quintile get back about twice what they pay in to the system, while people in the top quintile get back only half.[18] Note, however, that the CBO choice of discount rate (3.3 percent real; chosen to match the projected rate on Social Security Trust Fund assets) implies that the cohort average for the present value ratio is somewhat below 100 percent (a point which is discussed further in the sensitivity analysis below).

One of the primary reasons for the disparity between the CBO results and conclusions from the academic literature is that the CBOLT approach allows for a more comprehensive

[18]The interesting pattern across cohorts–falling then rising net returns–is due to two competing trends. First, the changes in Normal Retirement Age (NRA) for the baby boom and beyond permanently lowers net returns relative to earlier cohorts, but increasing life expectancy eventually overwhelms that effect, and net returns begin to rise for cohorts born after the 1970s in particular.

concept of population subgroup. Figure 1 includes everyone in a birth cohort who survives through the age of 45.[19] The most significant difference with respect to most published papers on progressivity is the inclusion of people who receive benefits from the disability insurance (DI) program. Those individuals generally have low earnings (as described in the last section, based on empirical patterns) but obviously get benefits at younger ages and for longer periods. There are of course other possible subgroups one could study, and one would expect that differences in present value ratios would get smaller as those groups get more homogeneous (this possibility is also explored in the sensitivity analysis below).

The final set of decisions needed to generate Figure 1 (or any estimates of redistribution) involve the way in which individuals in the micro sample are classified into groups. The groupings here are by quintiles of total lifetime earnings measured on a household basis. Specifically, the sum of a given person's real earnings over their lifetime is the basic classifier if they remain single in all years. If they happen to be married in a given year (CBOLT is an annual model) then the earnings measure for that year is a function of their earnings plus their spouse's earnings (adjusted for economies of scale). Thus, one of the standard criticisms of some progressivity measures–low earning spouses who are entitled to auxiliary benefits should be classified along with their high earning spouses, rather than as low earners–is already addressed in the baseline CBO estimates (however, as with the other two input decisions for the baseline, the issue of classifier is explored further below).

---

[19]The age of death cutoff chosen for inclusion in the distributional analysis is somewhat arbitrary. The determining factor in CBOLT is the goal of comprehensively sampling the 1940s cohort. Since the CBOLT micro sample is representative of living individuals as of 1984, the age 45 cutoff guarantees that people born in 1940 who survive through age 45 will be statistically represented.

Sensitivity of Progressivity Estimates Part 1: Discount Rates, Subgroups, and Classifiers

The choices underlying the estimates in Figure 1–discount rate used to measure net return, population subgroup studied, and earnings classifier selected–all have possible implications for the estimated redistribution in Social Security. More importantly, from the perspective of this paper, there is also a potential interaction between the estimated effect of differential mortality and each of the choices. That is, varying the rate of differential mortality may not have much impact given a particular set of choices for the discount rate, population subgroup, and classifier, but it might for some other choices for those three inputs. Thus, the first step in evaluating the effect of differential mortality on progressivity is to consider the overall sensitivity with respect to the three input choices and the possible interactions of each with differential mortality.

The first decision analyzed is the choice of discount rate used in the net return calculations (see Figure 5). One would expect that lowering the discount rate from the 3.3 percent real to some arbitrarily lower number (in this case, 2.0 percent real) would increase the present value of benefits relative to the present value of taxes because taxes are paid before benefits are received. Figure 5 shows that indeed is the case. The overall cohort present value ratios using a 2.0 percent discount rate are slightly above 100 percent, which suggests that 2.0 percent real is a close approximation to the cohort-level internal rates of return in the model. Although this choice does have a big impact on the level of present-value ratios (the lowest quintile gets back 2.5 to 3 times what they put in; the highest quintile something like 90 percent) the effect is similar across the quintiles, so the impact on progressivity estimates is much more modest.

There is reason to believe that the choice of discount rate could bias the inferences about the effects of differential mortality, however, so the alternative 2.0 percent real discount rate is used for the remaining calculations in this paper. The explanation goes like this: if one uses a discount rate that is too high, then increasing the extent of differential mortality (the gap between low and high earner lifespans) will not have a big impact on estimates of progressivity because the extra years of benefits that are shifted from low to high earners are discounted by an arbitrarily high factor.[20]

The second important input to consider is the choice of population subgroup (see Figure 6). One would expect that as the population subgroup becomes more homogeneous, the estimated progressivity would fall, and that is indeed the case. The progression of groups considered in Figure 6 is from the base case–the entire population surviving through age 45–to that group less the DI beneficiary population, to just the subgroup who qualifies for and lives long enough to receive Old Age Insurance (OAI) worker benefits.[21] Unlike the effect of changing the discount rate, the impact on progressivity of moving to increasingly more homogeneous populations is very significant. The present value ratios in the bottom quintile of lifetime earners fall by a factor of nearly two, while the ratios for the top quintile are affected much less. This result is intuitive: more of the people in the top quintile are OAI worker beneficiaries than in the bottom quintile, so the restriction has different effects. There is no obvious choice of the "correct" subgroup for the analysis of differential mortality below, so the remaining tables are all constructed to reflect outcomes across the possible choices.

---

[20]Liebman (2002) provides a nice discussion of the role of the discount rate in progressivity studies.

[21]The smallest group–OAI worker beneficiaries–is generally the one considered in studies that use example or representative earners, such as Garrett (1995), Steuerle and Bakija (1994), and Armour and Pitts (2004).

The third choice made in order to develop estimates of progressivity is the classifier itself (see Table 1). The base case classifier underlying the results in Figures 4 through 6 is total lifetime earnings on a household basis, as described above. There are a number of other viable approaches to classifying individuals, and Table 1 considers three alternatives for the four population subgroups in the 1960s birth cohort. The first alternative, which causes the system to look more progressive than in the base case, is individual lifetime earnings. The second and third, which cause the system to look slightly less progressive than in the base case, are lifetime earnings (on a household basis) through age 45 only, and potential earnings (again on a household basis) for individuals at age 45.

Classifying people in the sample by individual (rather than household) lifetime earnings has a strong effect on progressivity long recognized in the literature. Many of the people classified as low lifetime earners are spouses of high lifetime earners; indeed, they are able to remain out of the labor force because they have sufficient household-level resources to do so. What that implies, however, is that their estimated net present value ratios will be quite high, because they pay in very little and receive extensive auxiliary (spouse and survivor) benefits. The effect on progressivity of switching to the individual lifetime earnings classifier holds across all four population subgroups considered: the entire population, the DI beneficiary population, the non-DI beneficiary population, and the OAI worker beneficiary population. The effect is weakest for the OAI worker beneficiaries, which makes sense because there are (by definition) no spouse beneficiaries in that group (though there are OAI dual beneficiaries, who still benefit from a spouse's higher earnings).

The other alternatives to the base case classifier are intended to capture two other features of real world populations. The first uses the same lifetime earnings measure as in the base case, but the summation is truncated at age 45. That is, any earnings after age 45 have no impact on the individual's quintile position. This is meant to eliminate the possible reverse causation in the tabulations–people who die shortly after age 45 will have low earnings based on the total lifetime measure, but there is no reason (other than differential mortality) to suspect they will have low earnings based on the measure of earnings through age 45. The last classifier is intended to capture the effect of labor market choices and shocks. This measure is the "potential" household earnings for each person in the sample, defined as the earnings which would have occurred at age 45 if the individual (and spouse if one is present) had worked full-time and experienced no labor market shocks.

Truncating the summation of lifetime earnings at age 45 (the third classifier in Table 1) has relatively little impact relative to using total lifetime earnings (the second and base case classifier) but the explanation for that varies depending on the population subgroup. The primary reason is that lifetime earnings through age 45 are a good predictor of total lifetime earnings for most people, and for those people who experience a shock that breaks the connection between pre-45 and post-45 earnings, there are offsetting effects. In the total population, a person who earns enough while young to win a spot in the highest quintile but who then experiences a negative shock is more likely to get DI benefits but also more likely to die younger, which are offsetting events for net returns. In the OAI worker sample, the effect of using the lifetime earnings through age 45 classifier is even more modest, but that is because receipt of OAI benefits is itself an additional qualifier–one has to survive through age 62 at least to get OAI benefits, so all that is being tested is whether pre-45 lifetime earnings are a good predictor of total lifetime earnings (which they are).

Using potential household earnings at age 45 to classify people makes the estimates of progressivity lower than under the other three groupings. This classifier is best understood after a little background on the micro earnings process in CBOLT. Every person in the CBOLT micro model gets a randomly assigned "earnings differential" the year they first enter the labor market. The earnings differential term captures the well-known observation that overall age-sex-cohort-education earnings profiles have low explanatory power for any given individual, but adding a highly correlated idiosyncratic error term greatly improves the fit. The differential is effectively a percent deviation between the individual earnings and the average for the age-sex-cohort-education group of which the individual is a member.[22]

Adding the earnings differential to the predicted profile average is the first step towards solving for actual earnings in CBOLT. The earnings differential interacted with the age-sex-cohort-education group average represents what the individual would earn if they worked full time, experienced no unemployment, and experienced no other random "transitory" shocks to earnings. That concept is exactly what the potential earnings classifier means–what the individual would earn working full-time with no shocks. It is measured on a household basis because the same calculation is carried out for the individual's spouse (but when the individual is 45, not when the spouse is 45) and the sum is (as with the second and third classifiers) adjusted for economies of scale in consumption.

Using potential household earnings at age 45 lowers estimated progressivity relative to

[22]In order to simulate micro samples of earnings in which the correlation between annual and lifetime earnings are realistic the individual permanent differentials actually evolve slowly over the lifetime as well. See Harris and Sabelhaus (2002).

the base case classifier, but by no means eliminates Social Security redistribution. For the entire population, the present value ratio in the bottom quintile falls significantly relative to the other classifiers, which makes sense because the potential earnings of some low earners (non-working spouses in particular) are much higher than actual earnings. However, the effect is much more modest for the other population subgroups. In particular, under the potential earnings classifier, the present value ratios for the bottom quintile remain about one and a half times the ratio in the top quintile. Thus, using the base assumptions about differential mortality, CBOLT shows significant progressivity even for the most narrow subgroup (OAI worker beneficiaries) using the broadest classifier considered (potential household income at age 45).

Sensitivity of Progressivity Estimates Part 2: Differential Mortality

The analysis of Social Security redistribution in the last section shows that a wide range of progressivity estimates can be derived depending on how one specifies the measure of redistribution, the population subgroup being studied, and the classifier used to divide the population into quintiles. But the key question of this paper is how differential mortality affects those estimates. In this section, the rate of (non-DI) differential mortality is systematically varied in order to answer that question. The key insight from these experiments is that introducing arbitrarily large disparities in death rates by earnings does not fundamentally change the estimates of progressivity, and thus differential mortality itself has only a modest impact on the gap between statutory and actual redistribution.

The method by which differentials in death rates are systematically altered for these experiments is to operate directly on the coefficients in the Lillard and Panis (1999) mortality equation. The range considered here is from no differential mortality (all coefficients set to zero) to a tripling of the coefficient on earnings in the mortality equation. The result of varying the coefficients is to generate varying degrees of differentials in death rates (see Figure 7). If there is no differential mortality, the relative mortality rates by earnings quintile are all one. The base case differentials (the darkest bars in Figure 7) indicate that individuals in the lowest quintile have 1.4 times the average death rate within an age-sex group, while people in the highest quintile have a death rate of 0.8 times the average. Figure 7 shows that the extreme case considered–a tripling of the coefficient on earnings in the differential mortality equation–generates a range for relative mortality of 2.7 for the lowest quintile to 0.4 for the highest quintile. In short, controlling for age and sex, people in the lowest earnings quintile are something like seven times more likely to die than people in the highest earnings quintile.

Systematically varying the amount of differential mortality has only a modest impact on estimates of progressivity (see Table 2). The estimates here are (as in Table 1) for the 1960s cohort and use the "household lifetime earnings through age 45" classifier to divide the various populations into quintiles. The estimates show (again, as in Table 1) that population subgroup has a first-order impact on estimated progressivity. But, within any given population group, the effect of varying the assumption about differential mortality is in the right direction (more differential in death rates lowers the gap between present value ratios across quintiles) but the effect is second-order.

4. Differential Mortality and Social Security System Finances

The second major research issue pursued in this study is whether socioeconomic differentials in death rates have a significant impact on Social Security system finances. The question of how differential mortality affects system finances is largely independent of the analysis of redistribution in the last section, because conclusions about progressivity depend to some extent on the other inputs to the calculations (measure of redistribution, population subgroup, and classifier). However, the method used to investigate the connection (systematically varying the socioeconomic differential and re-solving the model) is also useful for investigating the relationship between differential mortality and summary measures of the Social Security financing imbalance. The conclusion about system finances turns out to be very similar to the inferences about progressivity. Relative to other determinants of system finances, differential mortality has a fairly modest impact on projections of Social Security's financial imbalance.

There are a number of different ways to measure the gap between Social Security revenues and outlays. The first major decision that goes into choosing a summary measure is whether to focus on present-value measures that summarize an entire stream of inflows and outflows over some time period or to use a snapshot of the gap between revenues and outlays at some point in time. Given the choice of present value or single-year gap measure, one must also specify the time period to be considered in the calculations. The second major decision is how to normalize the chosen measure, in particular, whether to divide by Social Security taxable payroll (the preferred OCACT measure) or overall GDP (which CBO uses in its published tables).

The measures of system finances using base-case differential mortality in CBOLT suggest the Social Security system is probably insolvent in the long run (see Table 3). The measures shown include the 75 year present value imbalance as a percent of taxable payroll, the100 year present value imbalance as a percent of GDP, and the gap between revenues and outlays as a percent of GDP in 2050 and 2100. The base-case differential mortality runs (second row of each panel in Table 3, consistent with CBO 2004-A) imply a funding gap of 1.05 percent of taxable payroll overall the 75 year horizon, and a present value gap of 0.57 percent of GDP over a 100 year horizon. On an annual flow basis, the expected gap between revenues and outlays is 1.39 percent of GDP in 2050, rising to 2.11 percent of GDP by 2100.

The direction of the effect on system finances from varying differential mortality is exactly what one would expect. Holding overall death rates by age and sex constant in the simulations, moving from a situation where relative death rates are constant across earnings groups (no differential mortality) to cases where the socioeconomic differential is arbitrarily large will make the Social Security system more expensive. If higher income beneficiaries are more likely to survive, average benefits paid in each year will be higher.

It is striking to observe how insensitive those estimates actually are to varying the mortality differential, however. The rows of Table 3 vary non-DI differential mortality from none to triple the Lillard and Panis (1999) estimate. The dramatic differences in rates of differential mortality (Figure 7) lead to fairly modest changes in the estimates of system financial imbalance, on the order of the sensitivity analysis for several inputs to the projections (SSA, 2004). For example, the range for the 75 year actuarial balance (OCACT's preferred measure of financial imbalance) range from .83 percent of payroll if differential mortality is turned off to 1.39 percent of payroll when the model uses triple the Lillard and Panis (1999) coefficient.

5. Conclusions

The analysis here suggests that the oft-stated arguments about differential mortality and Social Security are probably somewhat overstated. The general belief is that Social Security is not nearly as progressive in practice–some would argue not progressive at all—as it appears if one simply considers the statutory tax and benefit formulas. The argument is that socioeconomic differences in mortality work counter to progressivity because lower lifetime earners are more likely to die younger and therefore receive benefits for fewer years. A possible corollary to the argument is that socioeconomic differentials also affect Social Security system finances, because if the higher earners survive longer, average benefits paid out are larger than if life expectancies are equal across the earnings distribution.

Although the arguments about how differential mortality affects progressivity and Social Security finances are true, the effects appear modest at best. The estimates of progressivity are clearly affected much more by the choices made to estimate redistribution: the parameters used to calculate net returns, the population subgroup being studied, and the classifier used to divide the population into quintile groups. Holding those constant, the effect of differential mortality is not very strong. Similarly, the impact on system finances seem fairly small, especially when one considers how other inputs to the projections affect summary measures of the Social Security financial shortfall. Together, these observations imply that arguments for changing Social Security because either (1) differential mortality reverses the intended redistribution anyway or (2) we should establish (or re-establish) real progressivity are both somewhat overblown. Social Security as currently constructed does redistribute income from lifetime high-earners to lifetime low-earners.

6. References

Armour, Brian S., and M. Melinda Pitts (2004). "Incorporating Insurance Rate Estimates and Differential Mortality Into The Net Marginal Social Security Tax Rate Calculation." *Public Finance Review*, 32(6): 588-609.

Bosworth, Barry, and Gary Burtless (2004). "Supply-Side Consequences of Social Security Reform: Impacts on Saving and Employment," CRR Working paper 2004-1. Center for Retirement Research at Boston College (January).

Bosworth, Barry, and Gary Burtless (2002). "The Effects of Social Security Reform on Saving, Investment, and the Level and Distribution of Economic Well-Being," CRR Working paper 2000-2. Center for Retirement Research at Boston College (January).

Brown, Jeffrey R., Jeffrey B. Liebman, and Joshua Pollett (2002). "Estimating Life Tables That Reflect Socioeconomic Differences in Mortality," in *The Distributional Aspects of Social Security Reform*, Martin Feldstein and Jeffrey B. Liebman, eds. Chicago: University of Chicago Press.

Cohen, Lee, Eugene Steuerle, and Adam Carrasso (2001). "Social Security Redistribution by Education, Race and Income: How Much and Why?" Paper presented at the Third Annual Joint Conference for the Retirement Research Consortium, Washington, D.C., May 17-18.

Congressional Budget Office (2004-A). *The Outlook for Social Security*. Washington, D.C. (June). Updated on the CBO web site (www.cbo.gov) March, 2005.

Congressional Budget Office (2004-B). *Long-Term Analysis of H.R. 3821, the Bipartisan Retirement Security Act of 2004*. Washington, D.C. (July).

Congressional Budget Office (2004-C). *Long-Term Analysis of Plan 2 of the President's Commission to Strengthen Social Security* . Washington, D.C. (July).

Congressional Budget Office (2004-D). *Long-Term Analysis of Diamond-Orszag Social Security Plan* . Washington, D.C. (December).

Coronado, Julia Lynn, Don Fullerton, and Thomas Glass (2002). "Long Run Effects of Social Security Reform Proposals on Lifetime Progressivity," in *The Distributional Aspects of Social Security Reform*, Martin Feldstein and Jeffrey B. Liebman, eds. Chicago: University of Chicago Press.

Diamond, P.A., and P.R. Orszag (2003). *Saving Social Security: A Balanced Approach*. Washington, DC: The Brookings Institution Press.

Elo I.T., and K.P. Smith (2003). "Trends in Educational Differences in Mortality in the Untied States." Mimeo. Presented at Annual Meeting of the Population Association of America, Minneapolis, MN, May 2003.

Feldstein, Martin, (2005). "Structural Reform of Social Security." National Bureau of Economic Research Working Paper #11098. (February)

Fullerton, Don, and Brent Mast (2005). *Income Redistribution From Social Security*, Washington, D.C. The American Enterprise Institute Press.

Garrett, Daniel M. (1995). "The Effects of Differential Mortality Rates on the Progressivity of Social Security." *Economic Inquiry* 23:457-475.

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

Harris, Amy Rehder, and John Sabelhaus (2003) "Projecting Longitudinal Earnings for Long-Term Policy Analysis," Technical Paper 2003-2. Congressional Budget Office (April).

Lauderdale, D.S. (2001). "Education and Survival: Birth Cohort, Period, and Age Effects." *Demography* 38:551-61.

Liebman, Jeffrey B. (2002). "Redistribution in the Current U.S. Social Security System," in *The Distributional Aspects of Social Security Reform*, Martin Feldstein and Jeffrey B. Liebman, eds. Chicago: University of Chicago Press.

Lillard L. and S. Panis. (1999). "Demographic Projections." Mimeo. (May 1999), available at http://www.mrrc.isr.umich.edu/conferences/cp/cp99_lillard_demo.pdf

O'Harra, Josh, and John Sabelhaus (2002) "Projecting Longitudinal Marriage Patterns for Long-Term Policy Analysis," Technical Paper 2002-2. Congressional Budget Office (October).

O'Harra, Josh, John Sabelhaus, and Michael Simpson (2004) "Overview of the Congressional Budget Office Long-Term (CBOLT) Policy Simulation Model," Technical Paper 2004-1. Congressional Budget Office (January).

Pappas, G., S. Queen, W. Hadden, and G. Fisher (1993). "The Increasing Disparity in Mortality between Socioeconomic Groups in the United States, 1960 and 1986." *New England Journal of Medicine* 329:103-09.

Perese, Kevin (2002) "Mate Matching for Micro-Simulation Models," Technical Paper 2002-3, Congressional Budget Office (November).

Preston, Samuel H., and Irma T. Elo. (1995). "Are Educational Differences in Adult Mortality Increasing in the United States?" *Journal of Aging and Health* 7:476-96.

Social Security Administration Board of Trustees (2004). “The 2004 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds” (April) www.ssa.gov/OACT/TR/TR04/tr04.pdf.

Smith, Karen, Eric Toder, and Howard Iams (2003/2004). “Lifetime Distributional Effects of Social Security Retirement Benefits.” *Social Security Bulletin*, 65:33-62.

Steuerle, C. Eugene, and Jon M. Bakija (1994). *Retooling Social Security for the 21st Century: Right and Wrong Approaches to Reform*. The Urban Institute Press, Washington, DC.

Figure 1
Sources of Differential Mortality in CBOLT

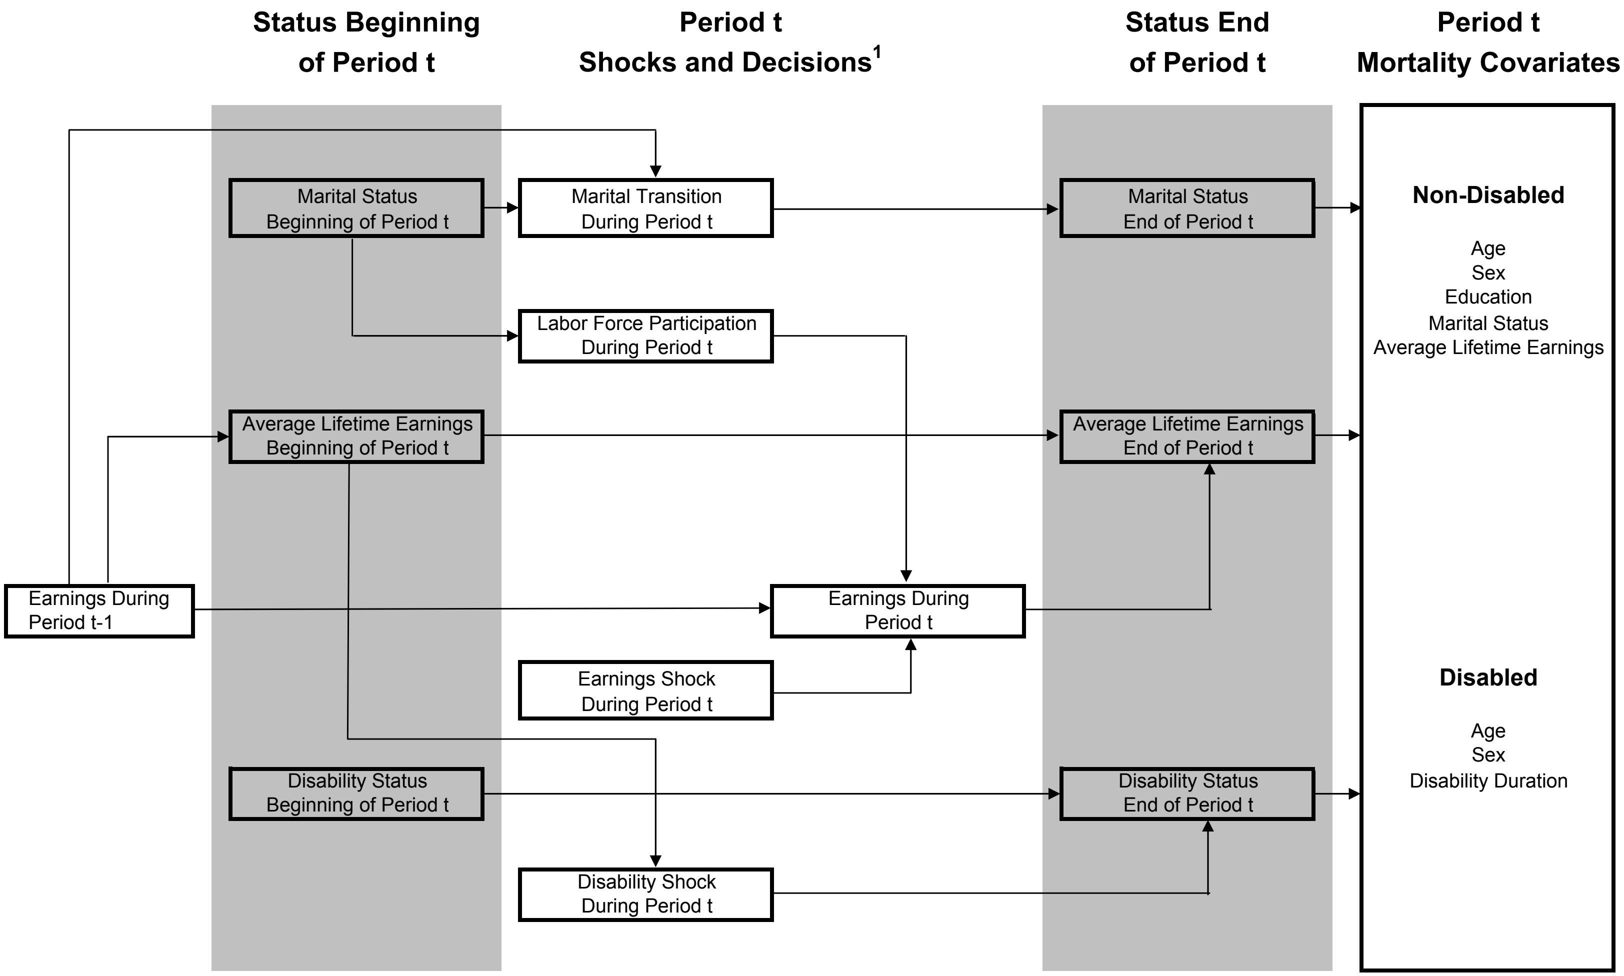

[1] In addition to the effects of the indicated inputs, all shocks and decisions are a function of age, sex, and education.

Figure 2. Historical Disability Insurance Incidence by Earnings Decile and Sex in the CWHS (Average 1984 to 2002, Ages 20 to 64)

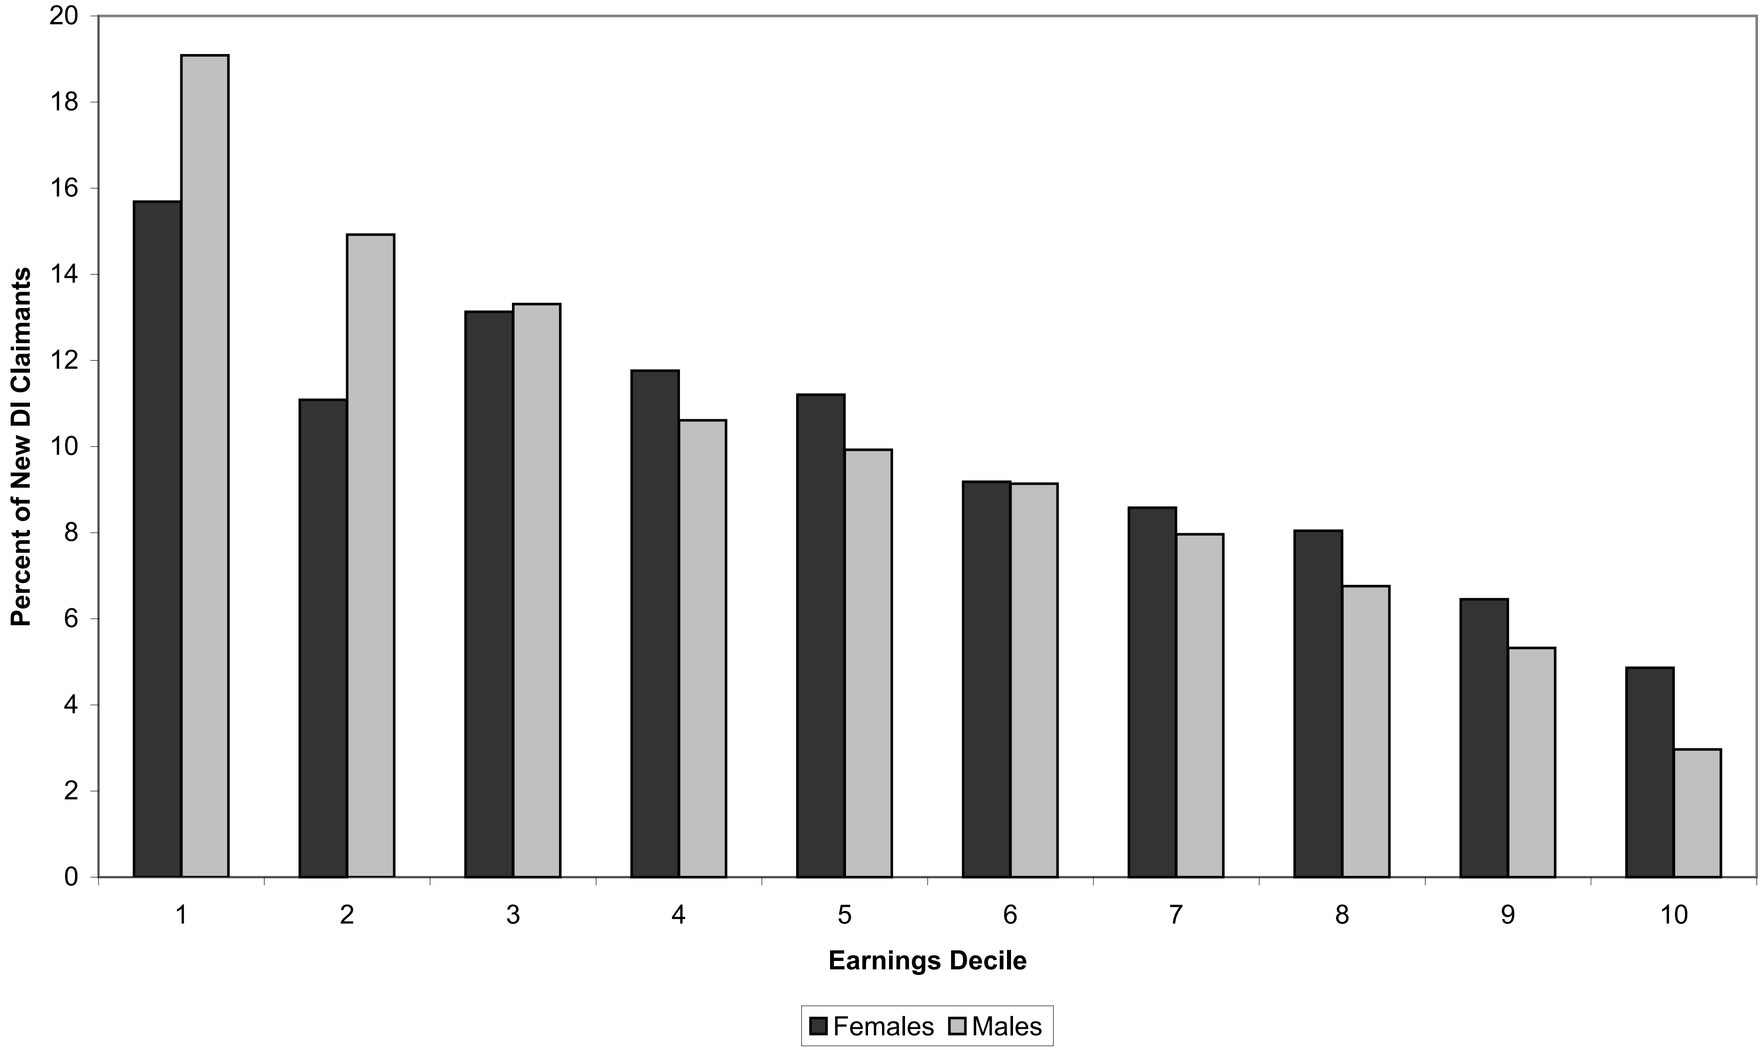

Figure 3. Social Security Administration Projected Death Rates
Disability-Insured (DI) Worker Beneficiaries by Duration
(Ages 30 to 54)

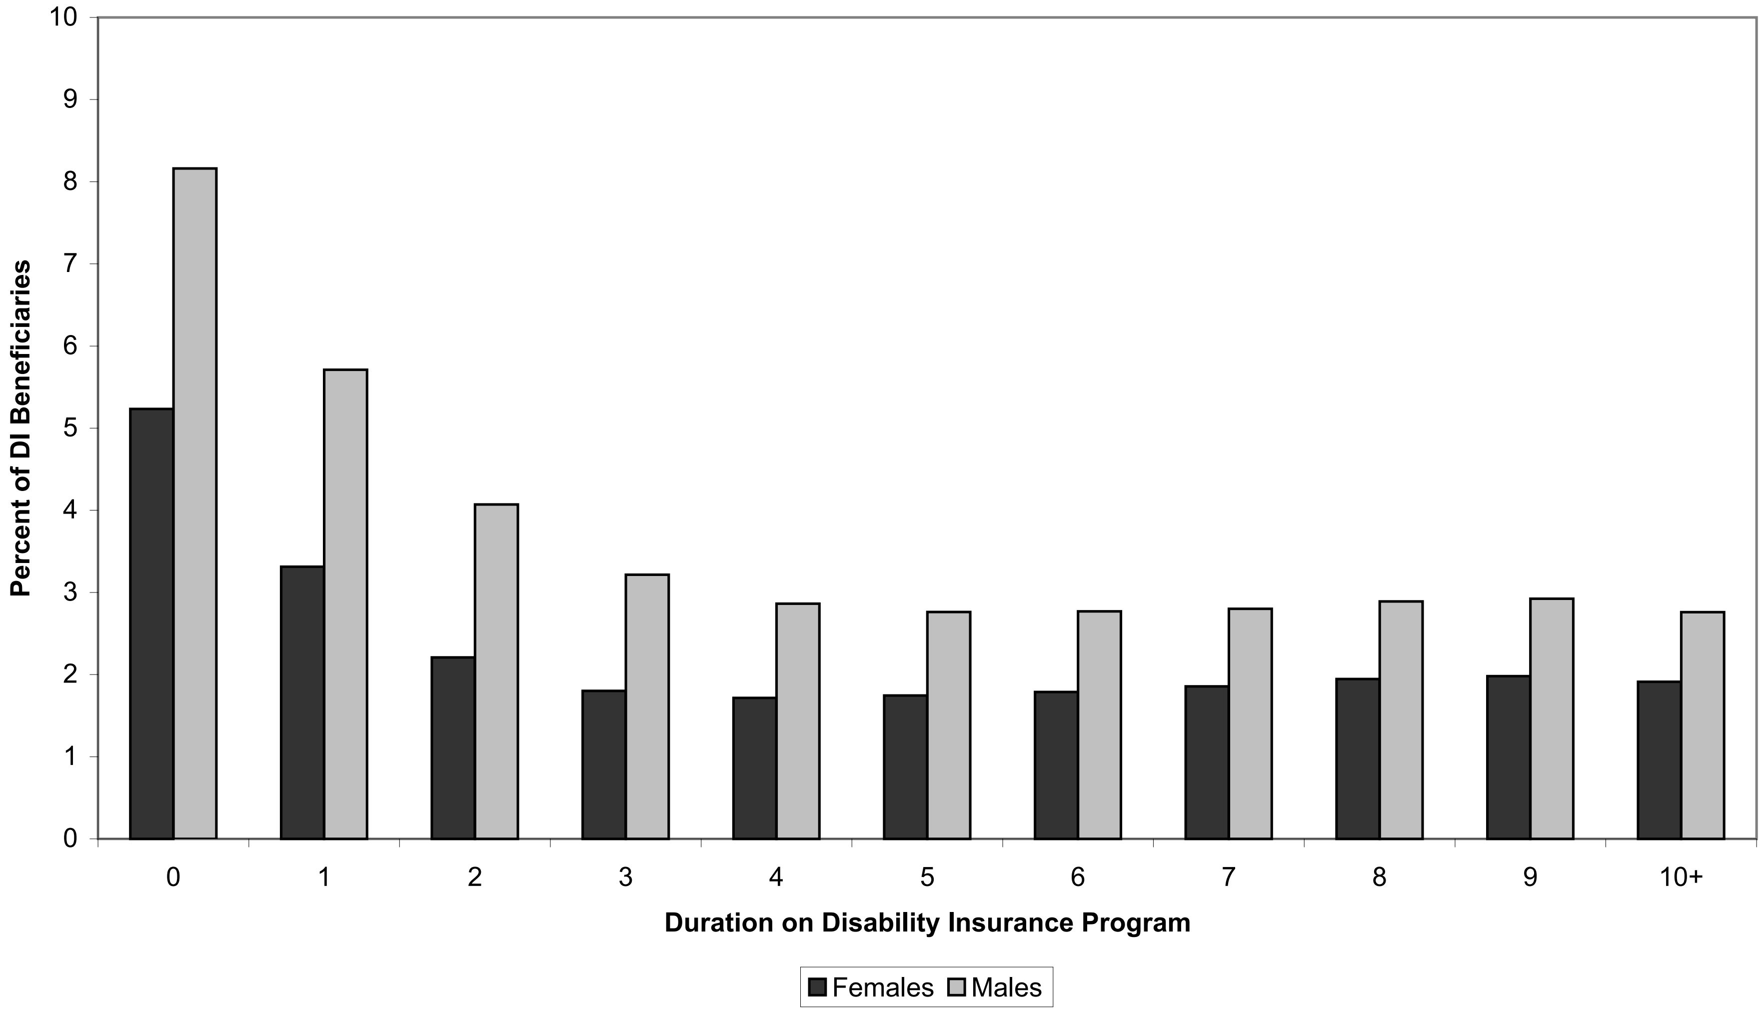

Figure 4. Projected Lifetime Benefit to Tax Ratios By Earnings and Cohort
(Total Population Surviving to Age 45, Classified by Household Lifetime Earnings, Discount Rate = 3.3%)

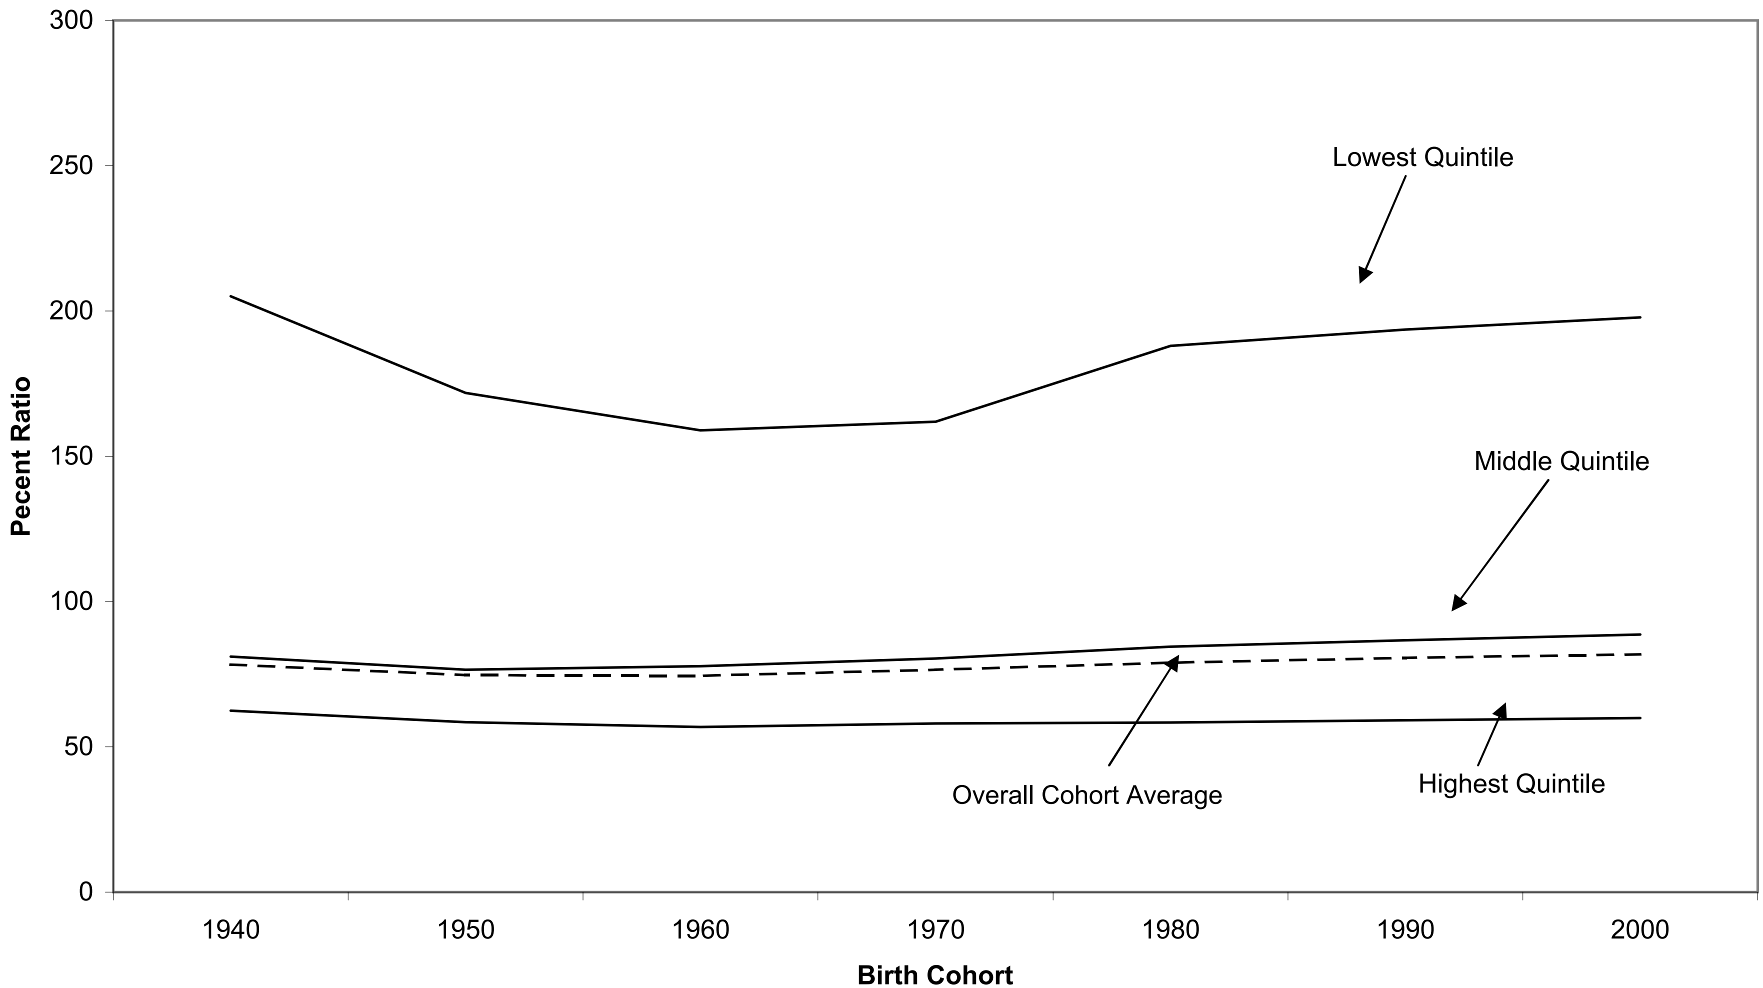

Figure 5. Effect of Discount Rate on Projected Lifetime Benefit to Tax Ratios
(Total Population Surviving to Age 45, Classified by Household Lifetime Earnings)

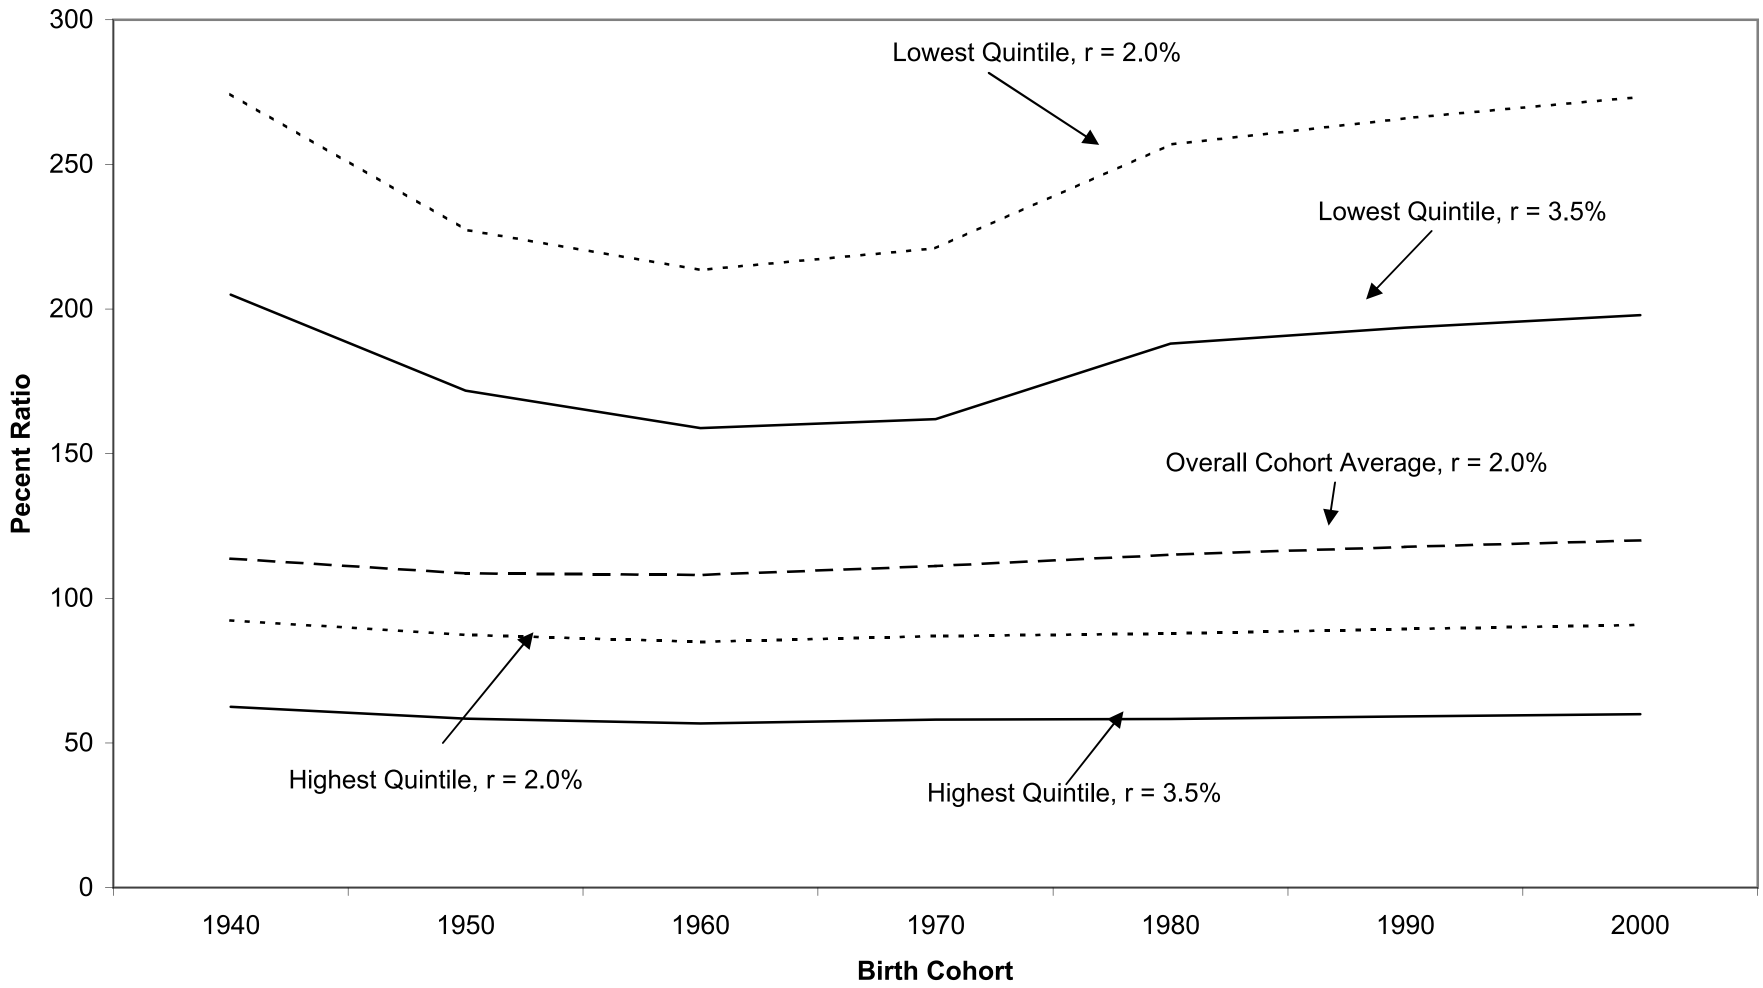

Figure 6. Effect of Population Subgroup on Benefit to Tax Ratios

(Various Populations Surviving to Age 45, Classified by Household Lifetime Earnings, Discount Rate = 2.0%)

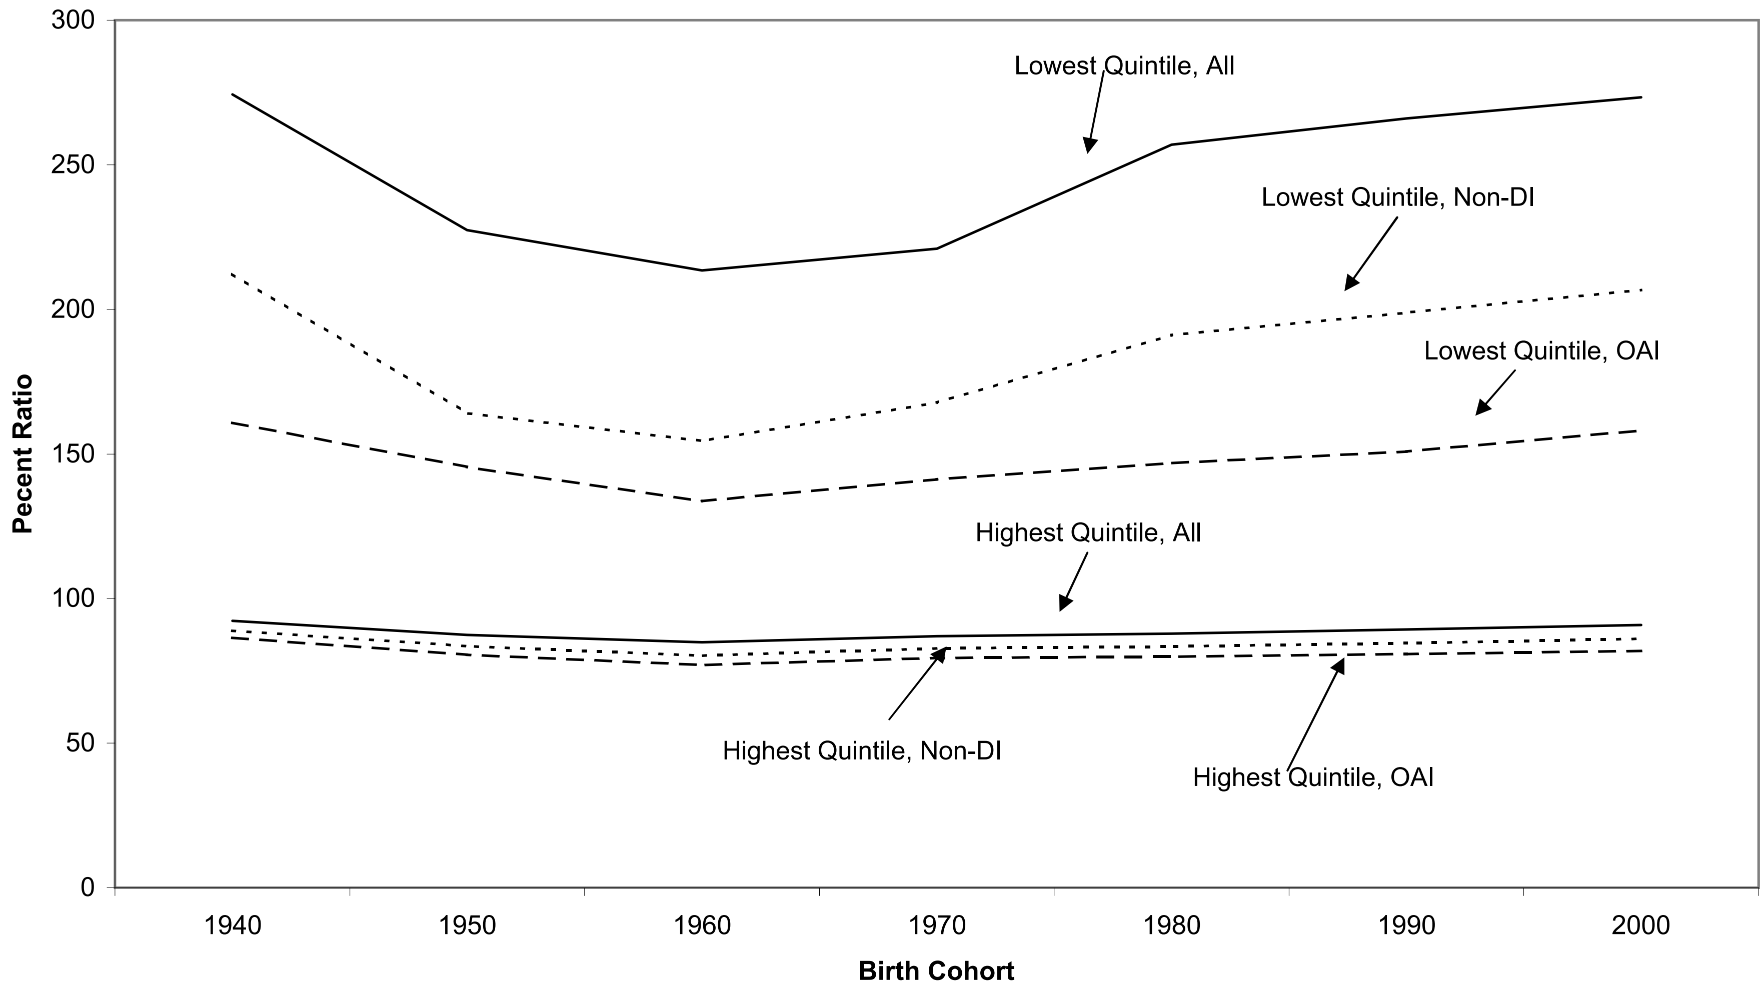

Table 1. Effect of Classifier on Projected Lifetime Benefit to Tax Ratios (1960s Cohort, Total Population Surviving to Age 45, Discount Rate = 2.0%)

	Percent Ratio of Lifetime Benefits to Lifetime Taxes Paid					
	Lowest Quintile	**Second Quintile**	**Middle Quintile**	**Fourth Quintile**	**Highest Quintile**	**Cohort Overall**
Entire Population, Classified by:						
Total Lifetime Individual Earnings	361%	170%	118%	96%	70%	108%
Total Lifetime Household Earnings	214	132	114	102	85	108
Total Household Earnings Through Age 45	182	137	113	100	87	108
Potential Household Earnings at Age 45	165	128	114	100	86	108
Disability Beneficiary Population, Classified by:						
Total Lifetime Individual Earnings	462	227	161	129	98	179
Total Lifetime Household Earnings	317	189	161	147	136	179
Total Household Earnings Through Age 45	315	213	172	153	141	179
Potential Household Earnings at Age 45	249	200	179	158	144	179
Non-Disability Population, Classified by:						
Total Lifetime Individual Earnings	305	148	107	89	68	95
Total Lifetime Household Earnings	154	111	102	94	80	95
Total Household Earnings Through Age 45	147	115	99	91	80	95
Potential Household Earnings at Age 45	139	110	101	89	78	95
Old Age Beneficiary Population, Classified by:						
Total Lifetime Individual Earnings	177	132	107	92	70	90
Total Lifetime Household Earnings	134	109	99	91	77	90
Total Household Earnings Through Age 45	130	109	96	88	77	90
Potential Household Earnings at Age 45	125	105	96	86	76	90

Figure 7. CBOLT Projected Relative Mortality Rates
(Mortality by Quintile Divided by Overall Mortality for Age/Sex Group)
(Males, 30 to 64, 2005 to 2110)

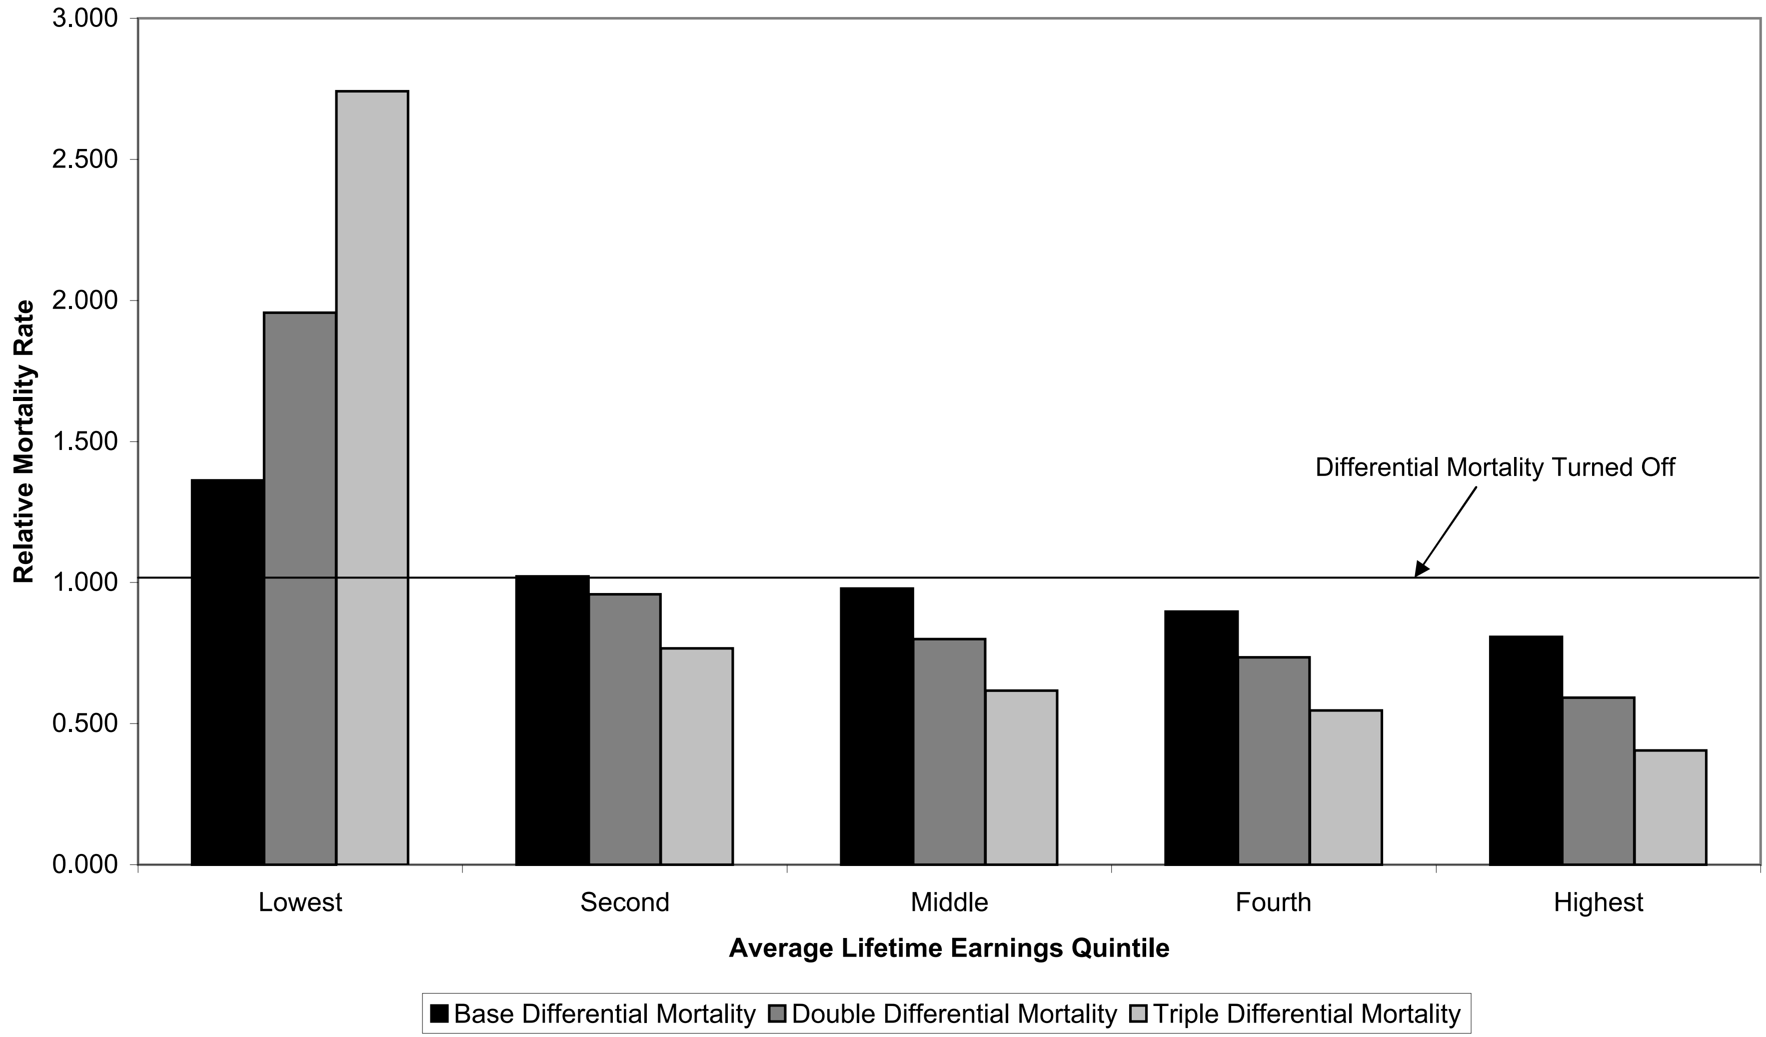

Table 2. Effect of Differential Mortality Assumption on Projected Lifetime Benefit to Tax Ratios (1960s Cohort, Classified by Total Household Earnings Through Age 45, Discount Rate = 2.0%)

	Percent Ratio of Lifetime Benefits to Lifetime Taxes Paid					
	Lowest Quintile	Second Quintile	Middle Quintile	Fourth Quintile	Highest Quintile	Cohort Overall
Entire Population, Differential Mortality:						
Non-DI Differential Mortality Turned Off	184%	138%	113%	98%	83%	106%
Base Case Differential Mortality	182	137	113	100	87	108
Non-DI Differential Mortality Doubled	178	136	114	102	91	110
Non-DI Differential Mortality Tripled	179	135	114	104	95	111
Disability Beneficiary Population, Differential Mortality:						
Non-DI Differential Mortality Turned Off	320	215	173	152	138	179
Base Case Differential Mortality	315	213	172	153	141	179
Non-DI Differential Mortality Doubled	311	211	171	154	144	179
Non-DI Differential Mortality Tripled	309	208	171	153	149	179
Non-Disability Population, Differential Mortality:						
Non-DI Differential Mortality Turned Off	150	116	98	88	76	93
Base Case Differential Mortality	147	115	99	91	80	95
Non-DI Differential Mortality Doubled	143	114	100	93	84	96
Non-DI Differential Mortality Tripled	144	114	101	95	87	98
Old Age Beneficiary Population, Differential Mortality:						
Non-DI Differential Mortality Turned Off	131	110	95	85	73	88
Base Case Differential Mortality	130	109	96	88	77	90
Non-DI Differential Mortality Doubled	128	108	96	89	80	91
Non-DI Differential Mortality Tripled	127	107	95	90	83	93

Table 3. Effect of Differential Mortality on Social Security System Finances

	Measures of OASDI Financial Status	
	Present Value Measures	
Differential Mortality Assumption	**75 Year Gap As A Percent of Payroll**	**100 Year Gap As A Percent of GDP**
Non-DI Differential Mortality Turned Off	-0.83	-0.48
Base Case Differential Mortality	-1.05	-0.57
Non-DI Differential Mortality Doubled	-1.21	-0.63
Non-DI Differential Mortality Tripled	-1.39	-0.70

	Gap Between Revenues and Outlays (Percent of GDP)	
Differential Mortality Assumption	**2050**	**2100**
Non-DI Differential Mortality Turned Off	-1.28	-2.00
Base Case Differential Mortality	-1.39	-2.11
Non-DI Differential Mortality Doubled	-1.48	-2.22
Non-DI Differential Mortality Tripled	-1.56	-2.31