For Whom the Bill Tolls:

Distributing the Costs of Maintaining

Social Security Solvency

John Sabelhaus*

May 2003

Abstract

Over the next several decades Social Security will most likely face a financial shortfall, and closing the funding gap will impose an incremental burden on some group at some point in time. One difficulty when measuring the distributional effects of proposed reforms is that risk characteristics vary across proposals. For example, adding an individual account component could introduce investment risk for participants, but it could also reduce the risk that benefit cuts or tax increases would be required. This paper distributes the costs of maintaining a solvent system under several proposed reforms using a Monte Carlo approach that reflects both expected outcomes and the various risks.

*U.S. Congressional Budget Office, Washington, DC 20515, johnsa@cbo.gov. The views expressed in this paper are those of the author and should not be interpreted as those of the Congressional Budget Office.

I. Introduction

At some point in the next few decades the U.S. Social Security system is likely to become insolvent if tax and benefit rules remain unchanged (Board of Trustees (2001), Congressional Budget Office (2001)). Reform proposals involving investment in private securities are sometimes characterized as a way of avoiding tax increases or benefit cuts, but when one considers the risks imposed on either beneficiaries or taxpayers, it is clear there is still an implicit cost that has to be allocated to someone (Smetters (2001), Feldstein, Rangulova, and Samwick (2001), Harris, Meyerson, and Smith (2001), Constantinides, Donaldson, and Mehra (2002)). The fact that increased risk represents a cost to those who bear that risk underscores the difficulty of comparing the distributional effects of proposed reforms. If risk characteristics differ, it is inappropriate to simply compare expected outcomes across reforms (Feldstein and Rangulova (2001)).

This paper distributes the costs of maintaining a solvent Social Security system under several proposed reforms using a Monte Carlo approach that captures expected outcomes and a wide range of risks. The strategy is to project Social Security tax and benefit outcomes across cohorts under several different policy scenarios, ranging from changes in current law tax and benefit parameters to more significant system overhaul involving individual accounts. Each policy scenario is simulated many times using Monte Carlo draws for the crucial economic and demographic inputs, so the results are probability distributions for outcomes, not point estimates. Further, each scenario is guaranteed to be solvent in every simulation because of a "fail-safe" clause that requires benefit cuts or payroll tax increases if the Social Security trust fund is in fact exhausted during the simulation. Thus, the approach captures the underlying risks in the current system as well as any changes in risks introduced by the reform being analyzed.

The various reforms considered lead to a wide range of outcomes across cohorts. The implication of maintaining current benefit policy until the Social Security trust funds run out of money is startling: if the default solution is to cut benefits when the trust fund is exhausted, retirees several decades from now could face benefit cuts on the order of fifty percent, while near term retirees would be held harmless. Some decrease in benefits for future retirees may be warranted because life expectancy is expected to rise by about five years in this century, but the Social Security normal retirement age (NRA) is only scheduled to rise two years. However, increasing the NRA by another three years would only marginally change system financial projections. Indeed, the projected outcomes for future retirees under the NRA increase would not change relative to a base case scenario where benefits are automatically reduced if the system becomes insolvent.

Systematically raising the NRA in line with life expectancy does affect conclusions about the long run system financial *trajectory*, however. The overall gap between taxes and benefits (relative to taxable payroll) is still negative but stabilizes after the baby boom dies off, which suggests that some level of trust fund balance would be consistent with long run solvency. That is, total inflows (taxes plus interest earned on the trust fund) would equal total outflows (benefits paid) so no further policy changes may be warranted. The key to this scenario is that some benefit reduction be put in place before the end of the baby boom's retirement, such that the trust fund is not depleted before the baby boom dies off. The model suggests this could be accomplished by a permanent reduction in new benefit awards of about 12.5 percent beginning in 2007 (when the first baby boomers reach age 62). If that policy were put in place, the baby boom would be handing off a system that is expected (given current long run economic and demographic projections) to remain basically solvent in perpetuity.

Reforms involving investment in private securities could improve the expected returns of some future retirees, though that conclusion varies across groups and proposals because of the implicit assumption about who is doing the prefunding. The two approaches considered here are investing fifty percent of the trust fund in equities and implementing a two percent of payroll individual account with a simple benefit offset mechanism. The trust fund equity investment policy does significantly improve expected rates of return (relative to a benefit cut base case) for some future retirees who might otherwise face shortfall-related benefit cuts, but the policy is likely to be overwhelmed by cost-income gap fundamentals in the very long run. The individual account approach improves expected rates of return for cohorts born after about 1970, and even the downside risk of that policy is negligible when compared to the risks associated with doing nothing. Cohorts born before 1970 may not prefer the individual accounts relative to other possible reforms, however, because they will (either explicitly or implicitly through the benefit cut mechanism) be funding the transition to individual accounts.

II. Distributing Risk and Return Across and Within Generations

The strategy of this paper is to simultaneously measure aggregate Social Security finances and person-level outcomes under various policy rules and (stochastic) states of the world. Estimating the probability distribution of individual outcomes requires three distinct systems working together: a Social Security budget model which solves for aggregate system financial outcomes based on key demographic and economic inputs, a macro-demographic model that generates those key inputs in a Monte Carlo setting, and a representative micro model which solves for person level outcomes that are consistent with the underlying stochastic environment and policy rules.

Social Security Budget Model

The role of the aggregate budget model in this investigation is to track Social Security finances over time given policy rules and values for the demographic and economic inputs. The aggregate model used here is intended to capture the basic features of budgetary projections made by the actuaries at the Social Security Administration (Frees, 1999). The crucial test of the budget model is that it exhibit responses to variation in policy parameters (tax rates, benefit formulas) and input assumptions (mortality improvement, fertility, immigration, wage growth, inflation, interest rates, unemployment, and disability rates) that match those in the actuaries' estimates. Given that, any set of alternative policy rules and/or stochastically generated input assumptions will lead to simulated changes in financial projections that basically match what the Social Security actuaries would predict.

The Social Security budget model has two independent sub-models, dealing separately with the demographic and financial aspects of the OASDI system. The demographic model tracks population using a "cell-based" approach, projecting person counts by single year of age, sex, and marital status groups. The inputs to the demographic model are annual fertility, total immigration, and rates of mortality improvement (by detailed age and sex). Demographic processes are calibrated to match Social Security "intermediate" projections when the inputs are set accordingly, and changes in the inputs (jointly or separately) produce changes in population counts that are also consistent with the sensitivity analysis reported annually by the SSA actuaries (for example, Board of Trustees, 2001).

The financial part of the Social Security budget model is also generally "cell-based" in nature.

The approach in the model is to track variables like labor force participation, average benefits, and beneficiary counts by age and sex, then multiply by the population counts from the demographic modules to solve for determinants of aggregate system financial flows. Although most values are averages within cells, some parts of the model have no distributional detail at all–for example, changes in taxable payroll are calculated using a log-normal approximation to the aggregate earnings distribution. However, some parts of the model rely on extensive micro detail–for example, the module for computing new OAI and DI worker benefit awards is based on a sample of about 14,000 actual new beneficiaries from 1996. Those 1996 beneficiaries are “aged” forward each year to represent beneficiary populations in future years, and benefits are calculated directly on the “aged” sample. The heterogeneity in the new beneficiary sample is needed to capture the interaction between earnings histories and policy rules.

The extent of behavioral response in the Social Security budget model is fairly limited. For example, labor force participation (within detailed age, sex, and marital status groups) is determined by equations estimated on pooled CPS data sets that capture the effects of major demographic and residual trends. There is some policy response in the participation equations, because the equations for age groups 62 through 70 include benefit replacement rates and earnings test variables on the right hand side.

Stochastic Macro-Demographic Model

The goal of the stochastic macro-demographic model is to generate values for the economic and demographic inputs to the Social Security budget model in a Monte Carlo setting. The demographic inputs are the rate of mortality improvement across detailed age and sex groups, the overall fertility rate, and the level of immigration. The economic inputs for the budget model are real wage growth, inflation, the unemployment rate, interest rates, and rates of disability incidence and termination. For simulations involving investment in private securities, the stochastic simulator also generates values for equity and corporate bond returns. The goal of the macro-demographic model is to create realistic stochastic variation in the annual values for each input, including correlations between those variables.

The stochastic macro-demographic model starts with Social Security Administration (SSA) intermediate projections to set central tendencies for each of the nine inputs. The SSA intermediate values are based on extrapolating historical averages for (generally) stationary processes like wage and price growth rates, and in that sense is consistent with the underlying time-series analysis used to build the macro-demographic model used here. Most of the stochastic inputs are treated as independent time-series processes, though inflation, unemployment, and interest rates are modeled together in a vector auto-regression (VAR) and short-run wage growth dynamics are affected by the other three economic variables. In all cases the specifications were chosen using standard time-series techniques (CBO (2001)).

The three demographic models are all estimated using data from the Social Security Administration (see Table 1). Mortality improvement is modeled as an AR(1) process for each of 42 separate age-sex groups. The data are available back to 1900. Although the equations are estimated separately, the rates of mortality improvement across age groups are correlated because the vector of innovations is drawn from a multi-variate normal distribution estimated using the historical error terms, which are correlated. The overall fertility process is characterized as an ARMA(4,1) model estimated on data back to 1917. Identifying a stationary characterization for the fertility process is somewhat complicated by distinct breaks in the series at various points in history, notably at the end of the baby boom. Experiments with an alternative to the ARMA representation (a first-difference model) did change implied fertility dynamics, but the effect on variability in system finances was modest (CBO, 2001). Finally, the immigration process is also dominated by distinct breaks in the time-series at various points in history, but in this case, because of changes in policy. Again, the ARMA(4,1) representation is most appropriate for immigration.

Three of the four economic variables (inflation, unemployment, and the real interest rate) passed the test for inclusion in a vector auto-regression (VAR) while the fourth (real wage growth) was rejected. The VAR model uses two annual lags for each of the three variables, and is estimated using data from BLS (unemployment and CPI-W inflation) and the Social Security Administration (real interest rates on new issues of OASDI trust fund assets) for the period 1954-1999. Unemployment rates are transformed using a log-odds ratio prior to estimation so the predicted values are constrained to the zero-one range in all cases. Although real wage growth does not lead the other three variables, it is strongly correlated with contemporaneous values. After controlling for current values of inflation, unemployment, and interest rates, the real wage growth residuals are white noise, which suggests the interesting dynamics occur through the other variables, because real wage growth (by itself) is highly auto-correlated.

The last two inputs to the Social Security budget model are disability incidence and termination.

Finding reasonable equations for variation in disability incidence and termination rates is made difficult because consistent data only exist back to 1975, and is further complicated because of changes in policy (stated and implicit) with respect to eligibility for the program. The problem is similar to the fertility model; there are clear breaks in the data that are (ex post) explainable, but it is not clear how to use that information when predicting future variability. Because there is no clear signal from the data, both models are specified as simple AR(1) processes.

In policy experiments involving trust fund investment in private securities or individual accounts the stochastic model generates values for equity and corporate bond returns using a random returns or "white noise" process. The alternative to white noise is to introduce some sort of mean reverting process, for example, one based on the dividend to price ratio. Annual stock yields are not correlated over time, but there is some evidence that the current level of the stock market is correlated with future returns, and therefore, future yields are, in a sense, "predictable" (see, for example, Campbell, Lo, and MacKinlay (1997) or Campbell and Shiller (1998)). The evidence of mean reversion is statistically weak, however, so (following Feldstein and Rangulova (2001)) the random returns process is adopted.[1]

---

[1]There are actually two implications of adopting a mean reverting process for equity returns. First, as Diamond (1999) observed, the starting point for stock market values matters for the expected long run return in order to keep market returns and underlying economic growth consistent with one another. Although recent declines in U.S. stock prices have ameliorated this observed divergence, the market is still (in terms of historical dividend price ratios) overvalued. The second implication of adopting mean reversion works in favor of equity-based solutions, however. Sabelhaus and Smith (2003) show that a mean reverting process with the same *annual* variability suggests much lower *cumulative* variability at frequencies of 10 years or more, because any divergence between underlying economic growth and stock returns is resolved through the mean reverting process.

Given the choice of a white noise process for equity returns, one still has to select values for expected returns and the standard deviation of returns in the simulations. Published data on equity yields are available from Ibbotson Associates (2001) back to 1926, and the exact time period used to measure the average and variance for yields matters a lot (see Table 2). The average yield for the entire time series (1926 through 2000) is 7.497 percent, but the period of rapid growth in the late 1990s has had a big impact, evidenced by the fact that the average for 1926 through 1994 is only 6.631 percent. It is not clear how the recent episode of stock market declines will end, but clearly at least some of the difference in those averages will be erased.

The standard deviation of stock yields in both time periods (1926 to 1994 and 1926 to 2000) is close to 20 percent. The last two rows of Table 2 show that the variability of yields has declined significantly since 1954. (Those lines are reported because the equation for corporate bond returns described below is estimated using data from 1954 forward). In the simulations involving equity investment, the model assumes a white noise process with expected returns of 6.5 percent (the value used by SSA actuaries) and a standard deviation of 19.795 percent, which should, if anything, bias the answers towards generating less positive gains from equity investments.

Although equity returns are an independent white noise process, real corporate bond yields are integrated with the rest of the macro model. The equation for corporate bond returns uses the three variables in the VAR; unemployment, inflation, and the real interest rate on government bonds. Thus the equation captures both the overall variability in corporate bond yields and the short-run macro-dynamics in the rest of the model. As noted above, although Ibbotson data on corporate bond yields is available back to 1926, the range for the other variables (in a consistent format) is limited to 1954 and later, so the equation is estimated for that period.

Example Worker Micro Model

The micro model is based on the same "example" workers used by the Social Security Administration in their distributional analysis of current law and alternative policies (SSA, 2001 and 2002). Outcomes for example workers are tied to the economic/demographic environment and the Social Security budget model, thus alternative policies (either specified changes in policy rules or invoked benefit cuts/tax increases because of financing constraints) will impact the example workers.

In every birth cohort there are four example workers, low, average, high, and maximum earners. The first three example workers have fixed age-earnings profiles which are derived from actual earnings patterns in Social Security administrative data files. Maximum earners are, by construction, assigned earnings equal to the taxable maximum in each year. In all cases, earnings for any given cohort at a given age are expressed relative to the average wage index, so earnings for example workers are tied to average wages (which is determined by the macro-demographic model) in every simulation.

The connection between the policy rules in the budget and micro models is also direct. Example workers pay taxes and earn benefits according to the rules in place during the year. Those rules include both specified policies (or policy alternatives) and any endogenous changes in policies (tax increases or benefit cuts) which result because of funding shortfalls. Example workers also get private accounts in the relevant alternative policy scenarios, with returns (depending on the specific portfolio assumption) generated by the stochastic simulator. Thus the micro model is consistent with the economic outcomes in the stochastic economic model, the rules used to make budgetary projections in the budget model, and interactions between the two through trust fund shortfall financing "fail-safes."

The remaining issue in the micro model is how to summarize person-level outcomes. The main criterion for a good yardstick is that it provide meaningful comparisons of overall policy outcomes, which suggests it should reflect both payroll taxes and benefits. Also, the goal is to quantify the variability in outcomes using the stochastic simulator, which suggests parsimony in choosing how many statistics to evaluate. One solution is to present the results in the form of internal rates of return (IRRs) across and within generations. The IRR is the value of the real rate of return which equates the present values of taxes and benefits. The calculation amounts to solving for the value of the internal rate of return (r) for which the following identity holds,

$$\sum_{s=0}^{64} t_s (1+r)^{100-s} = \sum_{s=65}^{100} b_s (1+r)^{100-s} p_s$$

where t_s and b_s are taxes paid and benefits received at age s, respectively, and B_s is the probability of surviving until age s conditional on surviving until age 65. IRRs do suffer some potential shortcomings when used as an overall welfare measure of how the Social Security system affects any given individual (Geanakoplos, Mitchell, and Zeldes, 1999), but they are still the most concise way to compare outcomes across time and income groups.

The IRRs presented here are constructed using an approach which is similar to the method developed by the Social Security actuaries (SSA, 2001). However, the differences in goals here suggest some important changes to the IRR calculations. For example, in the projections here, benefits can change after a person is claiming if the trust fund shortfall (fail-safe) rule calls for across-the-board benefit cuts. Therefore, the calculations here require explicit computations of benefits and survival probabilities after age 65, rather than just the calculation of a benefit award and expected life span conditional on reaching age 65 as in the SSA measures.

In addition to measuring (potential) benefit changes after age 65, there are other differences between the IRRs computed here and those reported by SSA, but the patterns of estimated returns across cohorts and income groups are still very similar. For example, the IRR estimates here do not include (probability-weighted) Disability Insurance (DI) benefits, but to be consistent, DI taxes are also ignored. The calculations here do not assign values for auxiliary (survivor, widow, spouse, and children's) benefits, but that omission biases only levels, not trends, in internal rates of return. On net, when compared to SSA, the estimates here show the same decreases in IRRs when moving from low to high earners at any point in time, the same increases in IRRs when switching from male to female survival probabilities, and the same (basic) declines in IRRs for future cohorts of retirees.

III. Risk and Return Under Current Law

Reforms involving investment in private securities are often characterized as adding risk to Social Security, as though there is no risk in the current system. In fact, there is a great deal of uncertainty about benefit levels in the current system for two reasons. First, there is uncertainty about economic determinants of benefits like wage growth and inflation, so even if rules are unchanged, benefit outcomes have some probability distribution. Second, there is uncertainty about how the rules will change if the trust fund is exhausted. The exact date of exhaustion in these simulations will vary with the (stochastic) demographic and economic environment, but it is very likely to occur in about forty years if current law is maintained. It is also the case that the trust fund may be exhausted under some of the alternative specifications for tax and benefit rules–just stating that a policy change is expected to lead to system solvency does not mean that solvency will occur in all states of the world.

Guaranteeing solvency in all simulations is important because an insolvent system will provide someone with an apparently higher rate of return due to the fact that the calculations are failing to assign some fraction of the costs of maintaining solvency. Therefore, most of the simulations below are run with a "fail-safe" clause in which either payroll taxes are raised or benefits reduced when the system approaches insolvency. Because the Social Security system is expected to become increasingly insolvent as far as the eye can see, imposing the fail-safe clause lowers projected internal rates of return for the future cohorts who will experience the tax increases or benefit cuts.

Figure 1 shows the probability distribution of internal rates of return for average earners in birth cohorts 1937 through 2000 in simulations which involve no fail-safe. The figure is derived by solving the model 1,000 times and computing the deciles of IRR outcomes for the average earner in each cohort. The estimates are produced using unisex cohort survival probabilities, so these IRRs are effectively averages across men and women. The middle line of the shaded area shows the median estimated IRR, which basically matches the trends projected by the Social Security Actuaries (SSA, 2001) in their analysis of example workers. The shaded area overall covers the range between the 10th and 90th percentiles. The uncertainty about IRRs for cohorts that are currently entering retirement reflects the fact that mortality is still uncertain, even though lifetime taxes and real benefits are known with certainty. Notice that the range of uncertainty expands for younger cohorts as other relevant determinants (like wage growth) start to affect the outcomes, but the overall range for the probability distribution stabilizes fairly quickly.

The insolvent system projections suggest an expected drop in the IRR for the average worker from about 2.6 percent for the 1937 cohort to about 2.3 percent for the 1960 cohort, reflecting the fact that the 1960 group will have paid taxes at a higher rate during their working lives and will receive less in benefits because of the scheduled increase in the normal retirement age. Social Security Administration estimates suggest that IRRs for cohorts born before 1937 were well above 2.6 percent–as much as a full percentage point higher for those born just two decades earlier. Interestingly, the insolvent system projections also suggest that IRRs would begin to rise again for cohorts born after about 1990, because the law changes are fully phased in, and reductions in mortality (which lead to higher lifetime benefits) start to increase lifetime net benefits. That increase in IRRs for the youngest cohorts under the insolvent system underscores why one should discount projections which do not have fail-safe provisions. Clearly, the same forces which are leading to higher IRRs (increased life expectancy) will also worsen system finances in the future, and someone will have to pay those costs.

In order to translate changes in IRRs into a measure of economic well-being, it is useful to consider how IRRs and benefit levels are related. Figure 2 plots IRRs for three example earners (low, average, and high) in the 1937 birth cohort. When benefits are at 100% of the scheduled value, the IRRs are 3.6%, 2.6%, and 2.1% for low, average, and high earners. If benefits are cut to 90% of the scheduled value, the IRRs fall to 3.2%, 2.3%, and 1.7%–a slope of about -.4. That is, a .4 percentage point reduction in the IRR is consistent with a 10% benefit cut. The relationship is clearly non-linear in the extreme (as benefits go to 0% of the scheduled amount the IRR goes to negative infinity) but Figure 2 shows it is fairly linear over most of the range considered here (benefit cuts up to about 50%). Applying this, one could say the expected long-run decrease in IRRs for average workers in the (implausible insolvent) case of about -.3 percentage points is equivalent to something like a 7 to 8 percent benefit cut relative to the 1937 cohort. Also, the long run range of IRRs between the 10th and 90th percentiles is plus or minus 0.5 percent, which suggests significant uncertainty about lifetime net returns even if no tax or benefit changes occur.

The projections of IRRs in the current law case with no fail-safe (Figure 1) are a straw-man because the likelihood that the OASDI system will remain solvent are negligible. Figures 3 and 4 show the projected deterioration in system finances using two complementary approaches. The first picture is the probability distribution for the gap between annual OASDI income and cost rates in each year 2002-2100, and the second is the distribution for the ratio of trust fund balances to expenditures (often called the "trust fund ratio") for the same time period.

The gap between income and cost rates in Figure 3 indicates two phases of deterioration in the relationship between system inflows and outflows; the first is a rapid deterioration associated with the baby boom retirement through about 2030, and the second is a slower decline associated with increased longevity. That second feature–continued deterioration in system fundamentals after the baby boom has died off–plays a prominent role in the discussion of reform scenarios in the next section. The graph also shows steadily growing uncertainty about the cost-income gap, because uncertainty about determinants like mortality improvement and fertility rates has cumulative effects.

The median projection for the cost-income gap in Figure 3 is generally consistent with SSA projections, though the gap is slightly more negative in every year because the model is solved using stochastic techniques, while the SSA projections are deterministic. The divergence occurs because symmetric variation in model inputs around expected values does not cause symmetric variation in model outputs–it is not simply a question of means versus medians. For example, even if it is equally likely that inflation is one percent above or below the expected value, it is not true that system finances respond symmetrically to movements in inflation above and below the expected values. Therefore, expected finances are slightly worse in the stochastic simulations than in a deterministic run with the same expected values for inputs (for details see CBO, 2001).

The expected deterioration in system finances is also evident in Figure 4, which shows the probability distribution for trust fund ratios (trust fund balances divided by expenditures). Although the exact year in which the trust fund will run out of money is uncertain, the expected year is around 2040 (as the youngest of the baby boomers approach age 80), and the fund is almost certain to run dry by around 2060. Figure 4 is consistent with the pattern of cost-income gaps in Figure 3, but shows the effect of (1) starting with a positive trust fund balance in 2002, and (2) enjoying several years in which receipts (taxes plus interest earned on the existing balance) are larger than expenditures (benefits paid). Again, the concept of solvency considered throughout this paper is that the trust fund never become exhausted, which means the ratio never crosses zero.

Because insolvency is not useful as a starting point for measuring risk and return, two solvent alternatives which start with current law tax and benefit rules are also considered. Figure 5 shows the median IRR estimates for average workers under current law with two types of fail-safe provisions: a payroll tax increase, and an across-the-board benefit cut. Both of the solvent projections show stark declines in IRRs for future cohorts of retirees relative to the median IRR for the insolvent system (reproduced from Figure 1) to levels (in the case of benefit cuts) below half what is projected for the 1937 cohort.

The expected decreases in IRRs for tax increase and benefit cut fail-safes differ in a predictable way across cohorts. Because only workers are affected by a tax increase, while everyone is affected by a benefit cut, the IRRs under a benefit cut are expected to fall sooner and more systematically. In a sense, the tax increase is between the unrealistic insolvent case where none of the shortfall is assigned, and the benefit cut case where all of the shortfall is assigned to cohorts who are alive. The tax cut case distributes the shortfall across both current and future workers, so some of the burden is imposed on cohorts not represented in the graphs. In that sense, the graphs (which already show 63 birth cohorts, requiring projections through 2100) are simply too short to capture the full impact of a tax increase fail-safe.

The estimates in Figure 5 for the insolvent and two fail-safe scenarios are median outcomes across many stochastic simulations. Each of the stochastic simulations have different levels of benefit cuts or tax increases because of differences in the economic and demographic environment. Figures 6 and 7 show the probability distributions for IRRs under the benefit cut fail-safe, and the probability distributions for the underlying benefit cuts. Unlike Figure 1, the range for IRR outcomes explodes for the youngest cohorts. The upside potential is not much different than Figure 1, because that represents the low probability states in which little or no benefit reduction is needed. The downside risk is alarming: IRRs approaching zero for the youngest cohort are plausible, which is (see Figure 2) equivalent to giving the 1937 cohort something like a fifty percent benefit cut.

The probability distribution of benefit cuts in the benefit cut fail safe with current law taxes and benefits is shown in Figure 7. The model triggers benefit cuts when the trust fund ratio falls below one–that is, the system could not pay for a year of benefits. Once invoked, the level of benefit cut is set to target a trust fund ratio of one from that point forward. Across the board benefit cuts could start to occur (with ten percent probability) as early as 2030; if that happens, the system finances are on such a negative trajectory that a large cut (about twenty percent) will be immediately necessary. If the demographic and economic situation is more favorable, the benefit cut will be delayed and reduced. Still, Figure 7 confirms that a large and rising benefit cut is likely if tax rates and benefit rules are unchanged.

Figures 8 and 9 show the same pair of graphs (IRR and policy change probability distributions) for the tax increase fail safe scenario. As with the medians, the probability distribution of the tax increase fail safe show less deterioration in IRRs for the cohorts considered, though the same exploding probability distribution for the youngest cohorts is evident. The probability distribution for tax increases is very similar to the benefit cut scenario, which makes sense, because they are both responses to the same deterioration in system finances. Either the tax increase and benefit cut fail-safe scenarios could serve as a benchmark against which to evaluate proposed policy changes in the next section. But as noted, the tax increase fail-safe imposes some of the costs of maintaining system solvency on unborn cohorts, and for that reason, the benefit cut seems a more appropriate starting point for thinking about alternatives. That is particularly true for reforms which focus on getting long-run Social Security cost and income rates moving parallel, rather than just increasing taxes as needed to cover ever-growing benefits. Also, the benefit cut fail-safe was used by the President’s Commission to Strengthen Social

Security in their approach to evaluating proposed individual account modifications (PCSSS, 2001). Therefore, in the next section, the benefit cut is used as a reference point for evaluating other sorts of proposed changes.

IV. Risk and Return Under Various Social Security Reform Proposals

The benefit cut fail-safe scenario introduced and analyzed in the previous section is an undesirable solution to Social Security's long-term financing problems for a number of reasons. A system which relies on across-the-board benefit cuts for retirees could impose significant hardships with little warning, and therefore making changes now (which gives people time to react) is inherently preferable. Also, reforming the system now offers the potential to develop a solution which is more generationally neutral–younger cohorts may not like the idea that the oldest baby boomers are (under current law) likely to get their scheduled benefits but then leave behind a system which is in disastrous financial condition. Finally, reform could involve investment in private securities, which offer an expected higher rate of return. Those higher returns compensate investors for investment risk, but in the current context, that is only one component of the overall risk faced by Social Security participants. Therefore, it is inappropriate to measure variable investment returns relative to scheduled benefits as though they will to be paid with certainty.

Increasing Normal Retirement Age in Line with Life Expectancy

The first reform considered is increasing the normal retirement age (NRA) in line with expected increases in longevity. Current law includes two phased increases in the NRA. The 1938 birth cohort, for example, has an NRA of 65 and two months, and that rises by two months per birth year until the 1948 cohort faces an NRA of 66. The NRA stays at 66 for several more years, then the 1955 cohort faces an NRA of 66 and two months, and again the NRA rises by two months per year until the 1960 and later cohorts face an NRA of 67. As noted, in the context of the micro model with fixed benefit claiming age of 65, this amounts to a benefit cut, because of the actuarial reduction associated with early retirement age.

Is an increase in the NRA of two years consistent with expected increases in longevity? Figure 10 suggests the answer is clearly no–life expectancy, conditional on reaching age 65, is projected to rise by about five years for both men and women in this century. Therefore, an increase in the NRA to age 70 is appropriate. In the reform simulations that increase is phased using the same pattern as is in place for the increases to ages 66 and 67. Every seventeen years the NRA is phased up (over a period of six years) by one more year, so the 1977 birth cohort will face an NRA of 68, the 1994 cohort an NRA of 69, and the 2011 birth cohort an NRA of 70.

The increase in the NRA would have a significant effect on long-run system finances, but probably not alter the underlying shortfall problem. Figure 11 shows the probability distribution for the cost-income gap under the NRA increase. Although the first phase of financial deterioration associated with the baby boom retirement is unaffected by the change, the long-run gap between cost and income is likely to become much more stable. This suggests optimism about long run solvency in the following sense: there exists some value for the trust fund ratio for which the system would be expected to remain solvent, so long as cost and income gaps are not diverging, because part of the annual costs are paid for by interest on pre-funded trust fund balances.

Although the stable cost-income gap suggests optimism, Figure 12 shows that current law is unlikely to generate a trust fund ratio which would meet the condition for perpetual solvency. Relative to current law (Figure 4) the probability distribution for trust fund exhaustion is little affected by the NRA change, though there is some improvement in the upside potential (90^{th} percentile). In retrospect this makes sense–the trust fund is likely to be exhausted before any of the additional NRA changes begin to affect finances. If the macro demographic environment is particularly favorable, then the trust fund balance may go a little further because cost rates are reduced, hence the improvement in the upside potential.

Increasing the NRA to age 70 would have a predictable effect on IRRs relative to current law. Figure 13 shows median IRRs under three scenarios, with and without the NRA increase. Raising the NRA in an insolvent simulation (no fail-safe) creates a predictable extension of the step-wise decline that is already in place for the age 66 and 67 increases under current law (Figure 5). The more interesting case is when the NRA increase is run with the benefit-cut fail safe operative. Figure 13 shows that the current law IRR (reproduced from Figure 5) and NRA increase IRR are about the same when the benefit cut fail-safe is operative, because the (median) necessary benefit cut is smaller than the cut associated with the NRA change, so the fail-safe is the operative change for most ranges of the probability distribution.

Although raising the NRA to age 70 is not in any sense a solution to the Social Security solvency problem, it is arguably part of any solution, at least in scenarios where future tax increases are not considered. Future generations may wish to tax themselves more and pay for longer retirements, but that is beyond the scope of this paper. Ultimately, the reason it makes sense to use the NRA increase as a benchmark is that the reforms discussed below are all characterized in terms of both fairness and sustainability. Both principles suggest that controlling for increased longevity should be a building block in any reform.

Reduce New Benefit Awards

As noted above, with an NRA increase to age 70 in place, there exists some trust fund ratio which would lead to perpetual solvency because interest earned on the fund would combine with taxes to cover benefits and real trust fund growth in every year. Note that implementing a reduction in new benefit awards at some point before the trust fund is exhausted would improve the chances of meeting any given trust fund ratio target, but it would also lower the necessary ratio needed for perpetual solvency, because the cost-income gap is less negative. Indeed, given the NRA increase, a new benefit award reduction of 12.5% beginning in 2007 (when the youngest baby boomers reach age 62) would (at the median) create this perpetually solvent scenario.

As with all the reforms to follow, the effects of the 12.5% benefit award reduction are shown in two ways. Figure 14 shows the median IRRs for average workers by birth cohort, and Table 3 lists 5 points on the probability distribution for IRRs (10^{th}, 30^{th}, 50^{th}, 70^{th}, and 90^{th} percentiles) for selected birth cohorts (1940, 1955, 1970, 1985, and 2000). In both Figure 14 and Table 3 one can see all of the reforms together, and compare them to the base case (increasing the NRA with a benefit cut fail-safe). All of these simulations use a benefit-cut fail safe.

The effect of a 12.5% new benefit award reduction beginning in 2007 is predictable. IRRs fall significantly for birth cohorts starting with 1945, because they experience the benefit cut, but are not likely to gain from the improved system solvency. Remember, the average person in those cohorts will (probably) have died before the trust fund (probably) runs out of money, so, other than for political economy reasons, they have nothing to gain from the benefit reduction. Later cohorts do have something to gain. Relative to the default case of doing nothing (Figure 14) the break even for the median outcome is somewhere around the 1967 birth cohort. The probability distribution (Table 3) shows that, although median outcomes for the youngest cohorts (1985 and 2000) are improved by the 12.5% benefit cut, the downside risk (10^{th} percentile) is just as bad. Why? In those cases, system finances have deteriorated so far that the 12.5% benefit reduction is not the operative margin, the benefit cut fail-safe has kicked in.

Trust Fund Equity Investment

Another option is to invest some of the trust fund in equities. Figure 14 and Table 3 show IRR outcomes for a 50% equity investment phased in by 2010. Unlike the 12.5% benefit cut, the trust fund investment strategy is not likely to impact near-term retirees adversely, and offers potential gains to all cohorts considered here. Indeed, all percentiles of IRR outcomes between 10^{th} and 90^{th} are improved. The results may of course be missing low probability but very adverse consequences, but within the range of likely outcomes, equity investment improves outcomes relative to just raising the normal retirement age.

Although everyone is better off under the trust fund equity investment reform, the gains are not equally distributed. Clearly, there is little gain through the middle of the baby boom (1955) because those cohorts have a good chance of (probably) dying before the trust fund (probably) runs out of money, even if there is no equity investment. The biggest winners are cohorts born in the late baby boom or just after, because they are the first group which is likely to experience benefit cuts if policy is not changed. The increase in expected returns effectively buys this group more time before benefit cuts are needed. Note that the gains are still positive for the youngest cohort (2000 birth year) but the overall pattern of declining IRRs in the default case is still evident.

Individual Accounts

The last set of reforms considered involves introducing “carved-out” private accounts into the system, in which taxpayers allocate a fraction of their payroll taxes to private accounts but agree to give up some benefits in return. The simulations here assume a 2% carve-out implemented for people under age 55 in 2005 along with a benefit offset computed using the opportunity cost to the trust fund. That is, the amount by which traditional benefits are reduced depends on how much one has diverted to the private account, accrued forward using the real rate of return on the OASDI trust fund. That hypothetical balance is annuitized (at the same real interest rate) at retirement, and the annuity value is the amount by which benefits are offset. The participant receives the actual account balance in the form of a fair annuity, which is assumed to be invested in a balanced portfolio of 50% equity, 30% corporate bonds, and 20% government bonds (which earn the trust fund interest rate).

An individual account carve-out like this has no expected present value cost to the Social Security system, and thus is the purest measure of the effect on individuals of accepting investment risk in exchange for potentially higher rates of return. Although the present value net cost is zero, there is an allocation of costs across generations that comes from the interaction of trust fund status, individual account funding, and the benefit cut fail safe. Funding accounts initially requires a decrease in income rates with no offsetting decrease in costs, because benefits are not immediately reduced. If the trust fund ratio falls below one during this transition, the benefit cut fail-safe kicks in and thus some of the transition cost is funded by current beneficiaries.

Figure 15 and Table 3 show the impact on IRRs from introducing individual accounts, with and without an underlying 12.5% benefit cut. In both cases, the median IRR for cohorts born after the 1970s is significantly improved by the individual account policy relative to only raising the NRA, investing the trust fund in equities, or cutting benefit awards by 12.5% in 2007. In addition to the positive effects on the median, the upside potential is much higher, and even the downside (10^{th} percentile) is improved relative to the other reforms.

The important distinction between the two individual account plans is who bears the cost of the transition. In the case where benefit awards are not reduced in 2007, there is a gradual reduction in IRRs, bottoming out at about 1.5% for the 1960 cohort. In the case where benefits are cut by 12.5% initially there is less chance that the benefit cut fail safe will be invoked during the transition, and thus the late baby boom and 1970s cohorts are not expected to bear as much of the burden–some of those costs have been shifted back onto the early baby boom retirees through the 12.5% benefit award cut.

V. Conclusions

The analysis here has described a number of different scenarios for how the Social Security insolvency problem might be addressed and the likely winners and losers across generations under various approaches. If no changes are made (or if one assumes the only change is that the normal retirement age will be raised in line with increased life expectancy), the outlook for future generations of retirees is bleak. Effectively, retirees in the next few decades will exhaust the trust fund and leave the system finances on a significant negative trajectory. Assuming that beneficiaries in the middle to late part of this century will be unable to roll the implicit debt forward to unborn generations, the implied benefit reduction is expected to be about fifty percent for those just born.

Alternative reform scenarios differ significantly in terms of winners and losers. The problem becomes one of political economy in the truest sense–near term retirees (the early baby boom) would lose under any reform, so they have no immediate incentive to agree to change. Late baby boomers will suffer some deterioration in net benefits if rules are not changed, but the reforms considered here would not significantly change the probability distribution of their returns, so it is not clear they have an immediate incentive to participate either. However, there are significant potential losses (relative to a reformed system) for birth cohorts after 1970, and that group will be paying the taxes during the baby boom's retirement. Will the post-1970s birth cohorts decide that the baby boom has rolled the implicit Social Security debt one generation too far? If so, baby boomers have a very good reason to participate in the reform debate now.

VI. References

1994-1996 Advisory Council on Social Security. *Report of the 1994-1996 Advisory Council on Social Security, Volume I: Findings and Recommendations*. Washington, D.C.: U.S. Government Printing Office, January 1997.

Board of Trustees, Federal Old-Age And Survivors Insurance and Disability Insurance Trust Funds. *The 2001 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds*, Washington, D.C.: U.S. Government Printing Office, March 2001.

Campbell, John Y., Andrew W. Lo, and Craig A. MacKinlay. *The Econometrics of Financial Markets.* Princeton: Princeton University Press, 1997.

Campbell, John Y. and Robert J. Shiller. "Valuation Ratios and the Long-Run Stock Market Outlook," *The Journal of Portfolio Management* 24, No. 2 (Winter, 1998): 11-26.

Congressional Budget Office. *Uncertainty in Social Security's Long Term Finances: A Stochastic Analysis*. Washington, D.C.: Congressional Budget Office, December 2001.

Constantinides, George M., John B. Donaldson, and Rajnish Mehra. "Junior Must Pay: Pricing the Implicit Put in Privatizing Social Security." National Bureau of Economic Research Working Paper 8906, April 2002.

Diamond, Peter A. "What Stock Market Returns to Expect for the Future?" Center for Retirement Research at Boston College, *Issue in Brief*, No. 2. September, 1999.

Feldstein, Martin, Elena Rangulova, and Andrew Samwick. "The Transition to Investment Based Social Security when Portfolio Returns and Capital Profitability are Uncertain," in *Risk Aspects of Social Security Reform*. Eds. John Campbell and Martin Feldstein, Chicago: University of Chicago Press, 2001.

Feldstein, Martin, and Elena Rangulova. "Individual Risk in an Investment Based Social Security System," *The American Economic Review* 91, No. 4 (September, 2001): 1116-1125.

Frees, Edward W. (Jed). "Summary of Social Security Administration Projections of the OASDI System." Working paper for the 1999 Technical Panel on Assumptions and Methods (December, 1999).

Geanakoplos, John, Olivia S. Mitchell, and Stephen P. Zeldes, "Social Security Money's Worth," in *Prospects of Social Security Reform*. Eds. Olivia S. Mitchell, Robert J. Meyers, and

Howard Young. Philadelphia: Pension Research Council, Wharton School of the University of Pennsylvania, 1999.

Harris, Amy Rehder, Noah Meyerson, and Joel Smith. "Social Insecurity? The Effects of Equity Investments on Social Security Finances," *National Tax Journal* 54 No.3 (September 2001): 645-668.

Harris, Amy Rehder, John Sabelhaus, and Michael Simpson. "OAI Benefit Uncertainty Under Individual Investment Alternatives," Unpublished Paper, Washington, D.C.: Congressional Budget Office, May 2002.

Ibbotson Associates. *Stocks, Bonds, Bill and Inflation 2001 Yearbook, Market Results for 1926-2000.* Chicago: Ibbotson Associates, 2001.

President's Commission to Strengthen Social Security (PCSSS). "Strengthening Social Security and Creating Wealth for All Americans," Report of the President's Commission, December, 2001. Also available on the Commission web site at http://www.csss.gov/reports/Final_report.pdf.

Sabelhaus, John, and Joel V. Smith. "Projecting Equity Returns for Long Run Policy Analysis." Unpublished Paper, Washington, D.C.: Congressional Budget Office, May, 2003.

Social Security Administration, Office of the Chief Actuary. "Internal Real Rates of Return Under the OASDI Program for Hypothetical Workers." *Actuarial Note* Number 144, June 2001.

Social Security Administration, Office of the Chief Actuary. "Estimates of Financial Effects for Three Models Developed by the President's Commission to Strengthen Social Security," Memorandum to Daniel Patrick Moynihan and Richard D. Parsons from Stephen C. Goss and Alice H. Wade, January 31, 2002.

Smetters, Kent. "The Effect of Pay-When-Needed Benefit Guarantees on the Impact of Social Security Privatization" in Risk Aspects of Social Security Reform. Eds. John Campbell and Martin Feldstein, Chicago: University of Chicago Press, 2001.

Table 1
Equations in Stochastic Macro-Demographic Model

Variable	*Description of Stochastic Process*
Mortality Improvement	Separate AR(1) equations for each of 21 age and 2 sex groups estimated using SSA data for 1900-1995. Model draws 2 sets of 21 correlated errors across using multi-variate normal distribution.
Fertility	ARMA(4,1) equation for overall fertility rate estimated using SSA data for 1917-1997.
Immigration	ARMA(4,1) equation for total immigration estimated using SSA data for1901-1995.
Unemployment	VAR model with two lags each on unemployment, inflation, and real interest rate estimated using BLS and SSA data for 1954-1999.
Inflation (CPI-W)	VAR model with two lags each on unemployment, inflation, and real interest rate estimated using BLS and SSA data for1954-1999.
Real Wage Growth	Level of nominal wage growth a function of the three economic variables, equation estimated using NIPA, BLS, and SSA data for 1954-1999; real wage (comparable to SSA's "differential") is nominal wage less inflation.
Real Interest Rate	VAR model with two lags each on unemployment, inflation, and real interest rate estimated using BLS and SSA data for1954-1999.
DI Incidence	AR(1) model for overall DI incidence rate estimated using SSA data for 1975 to 1999.
DI Termination	AR(1) model for overall DI termination rate estimated using SSA data for 1975 to 1999.
Equity Returns	Total returns a white noise process estimated using Ibbotson Associates (2001) large-cap data for the period 1954-1999.
Corporate Bond Returns	Level of large-cap bond returns a function of inflation, unemployment, and interest rate estimated using Ibbotson Associates (2001), SSA, and BLS data for the period 1954-1999.

Table 2
Historical Equity Returns and Variability of Returns

Time Period	Mean	Standard Deviation
1926 to 2000	7.497	19.593
1926 to 1994	6.631	19.795
1954 to 2000	8.014	16.793
1954 to 1994	6.724	16.769

Source: Ibbotson Associates (2001)

Figure 1: Distribution of Internal Rates of Return
Current Law, No Fail-Safe (Insolvent System)

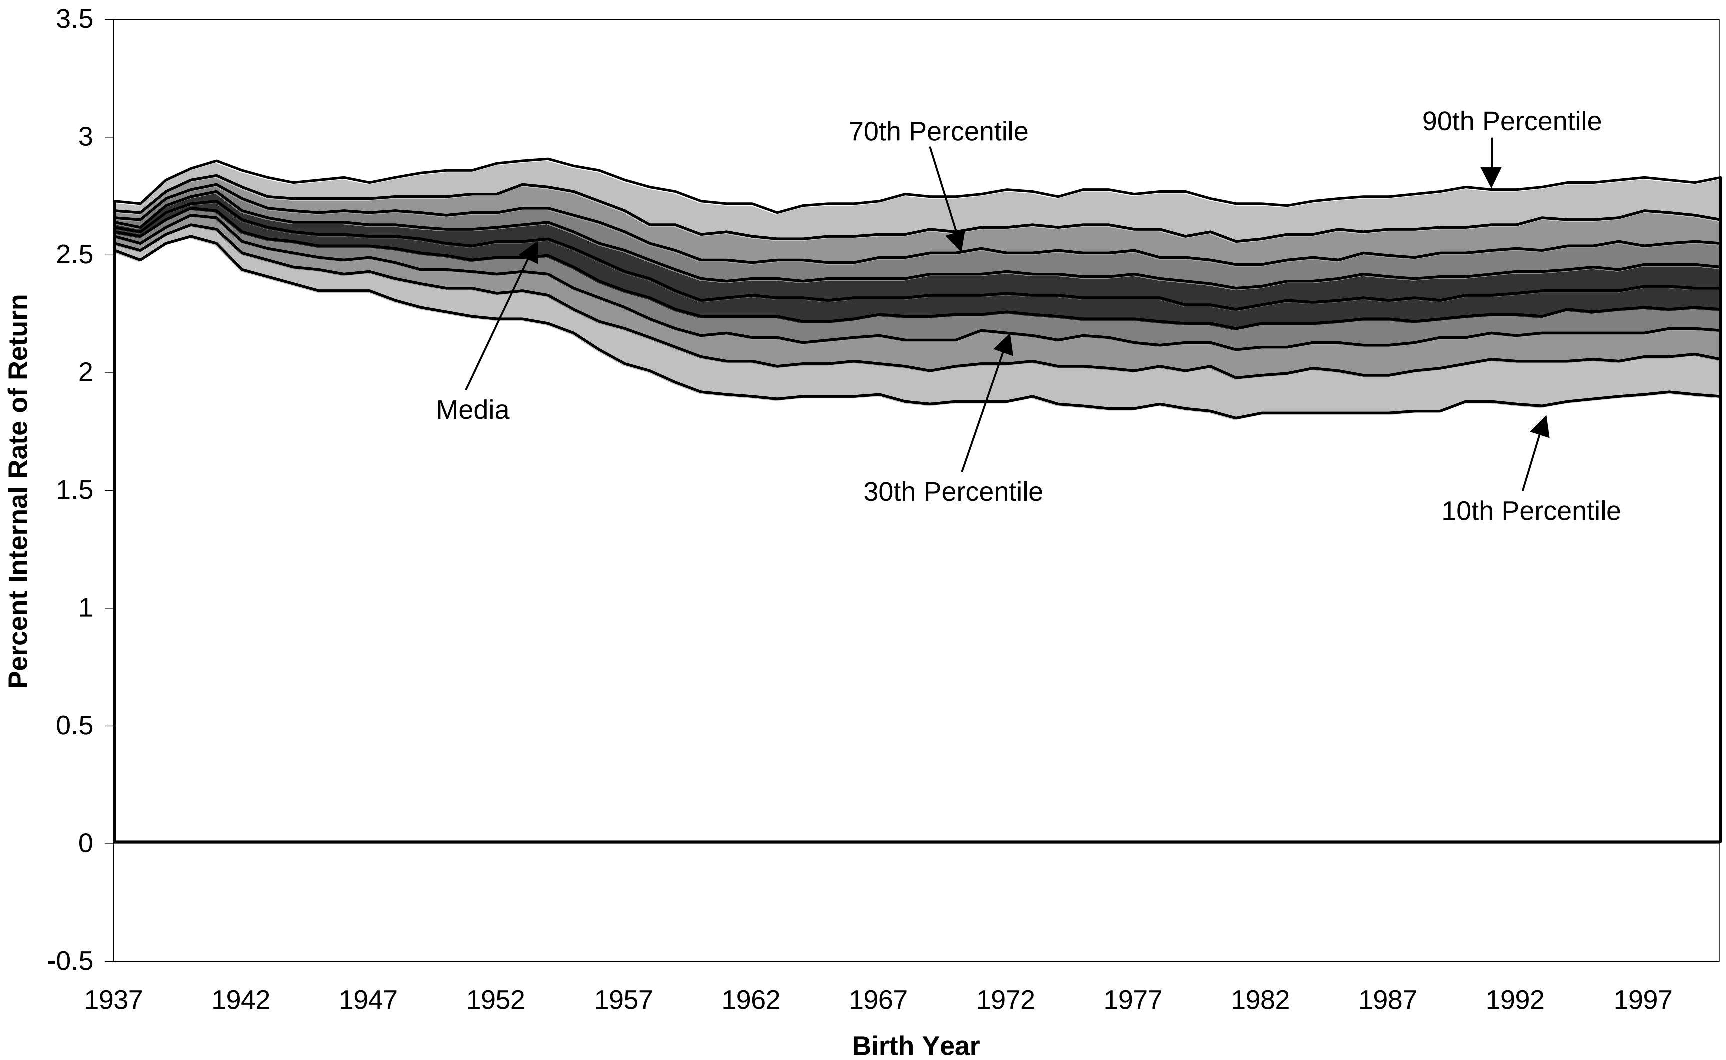

Figure 2: Benchmark Relationship Between IRRs and Benefit Cuts (1937 Birth Cohort)

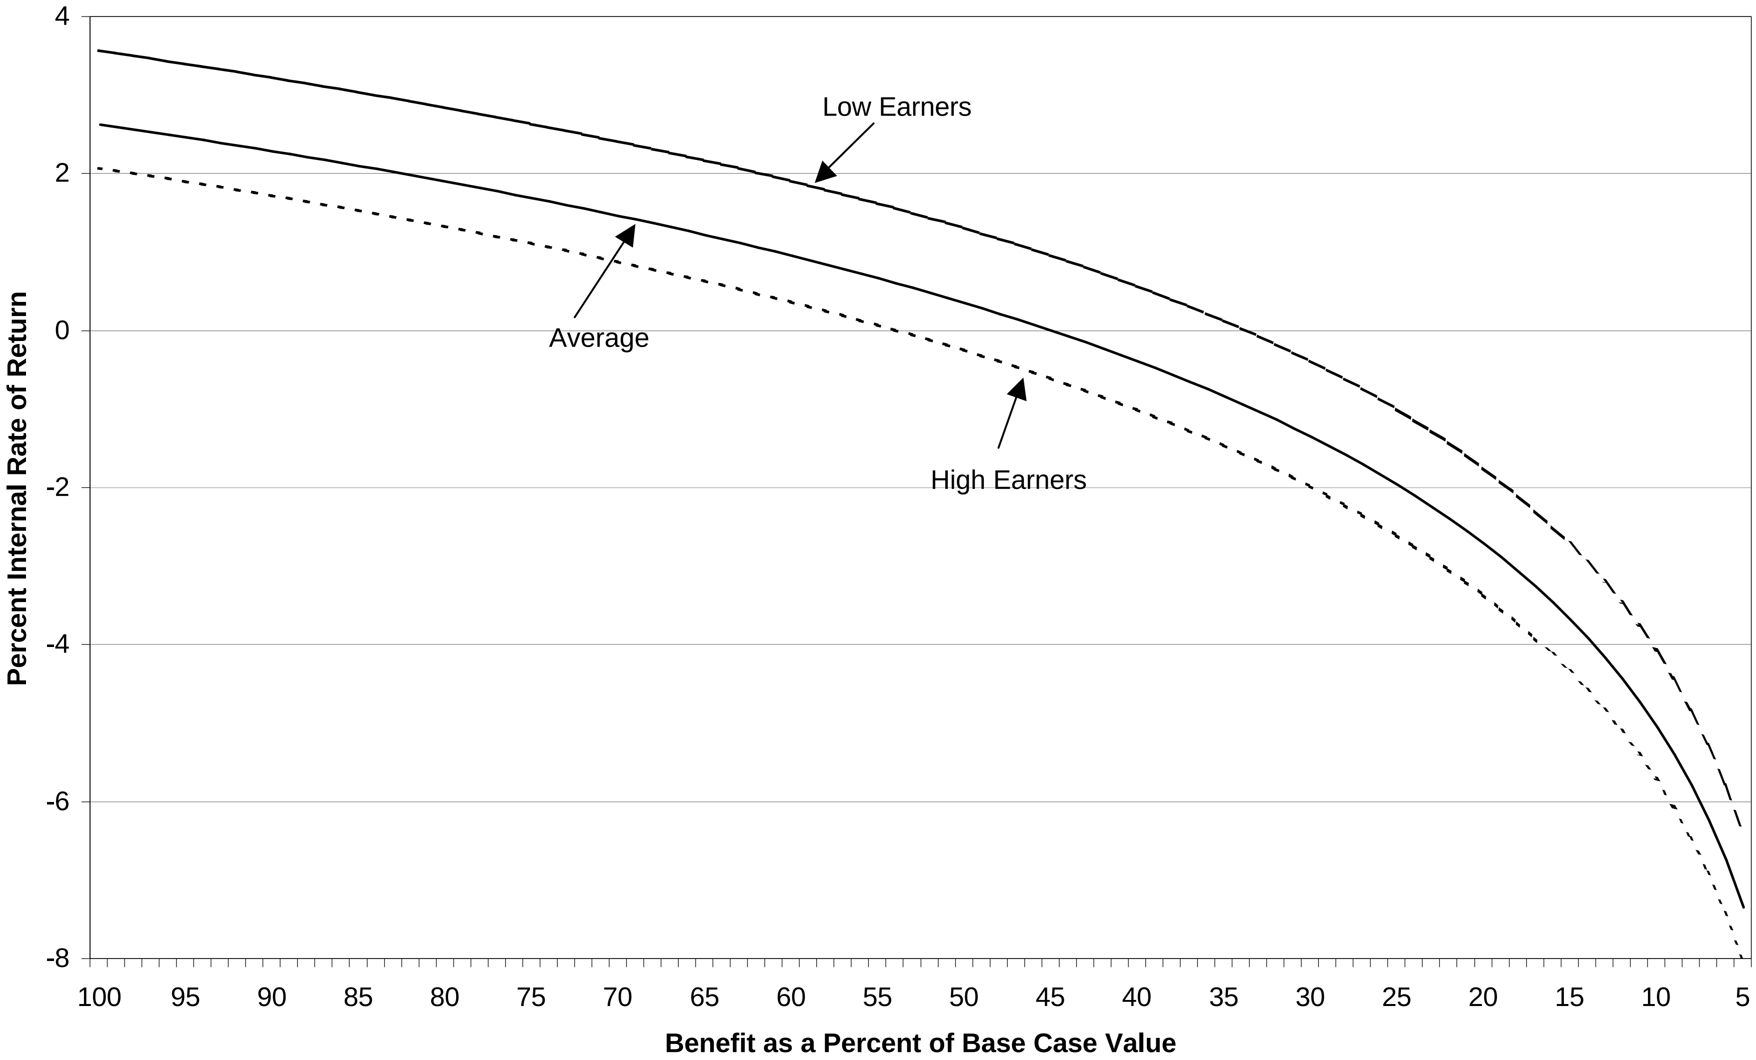

Figure 3: Distribution of Gap Between Cost and Income Rates
Current Law, No Fail-Safe (Insolvent System)

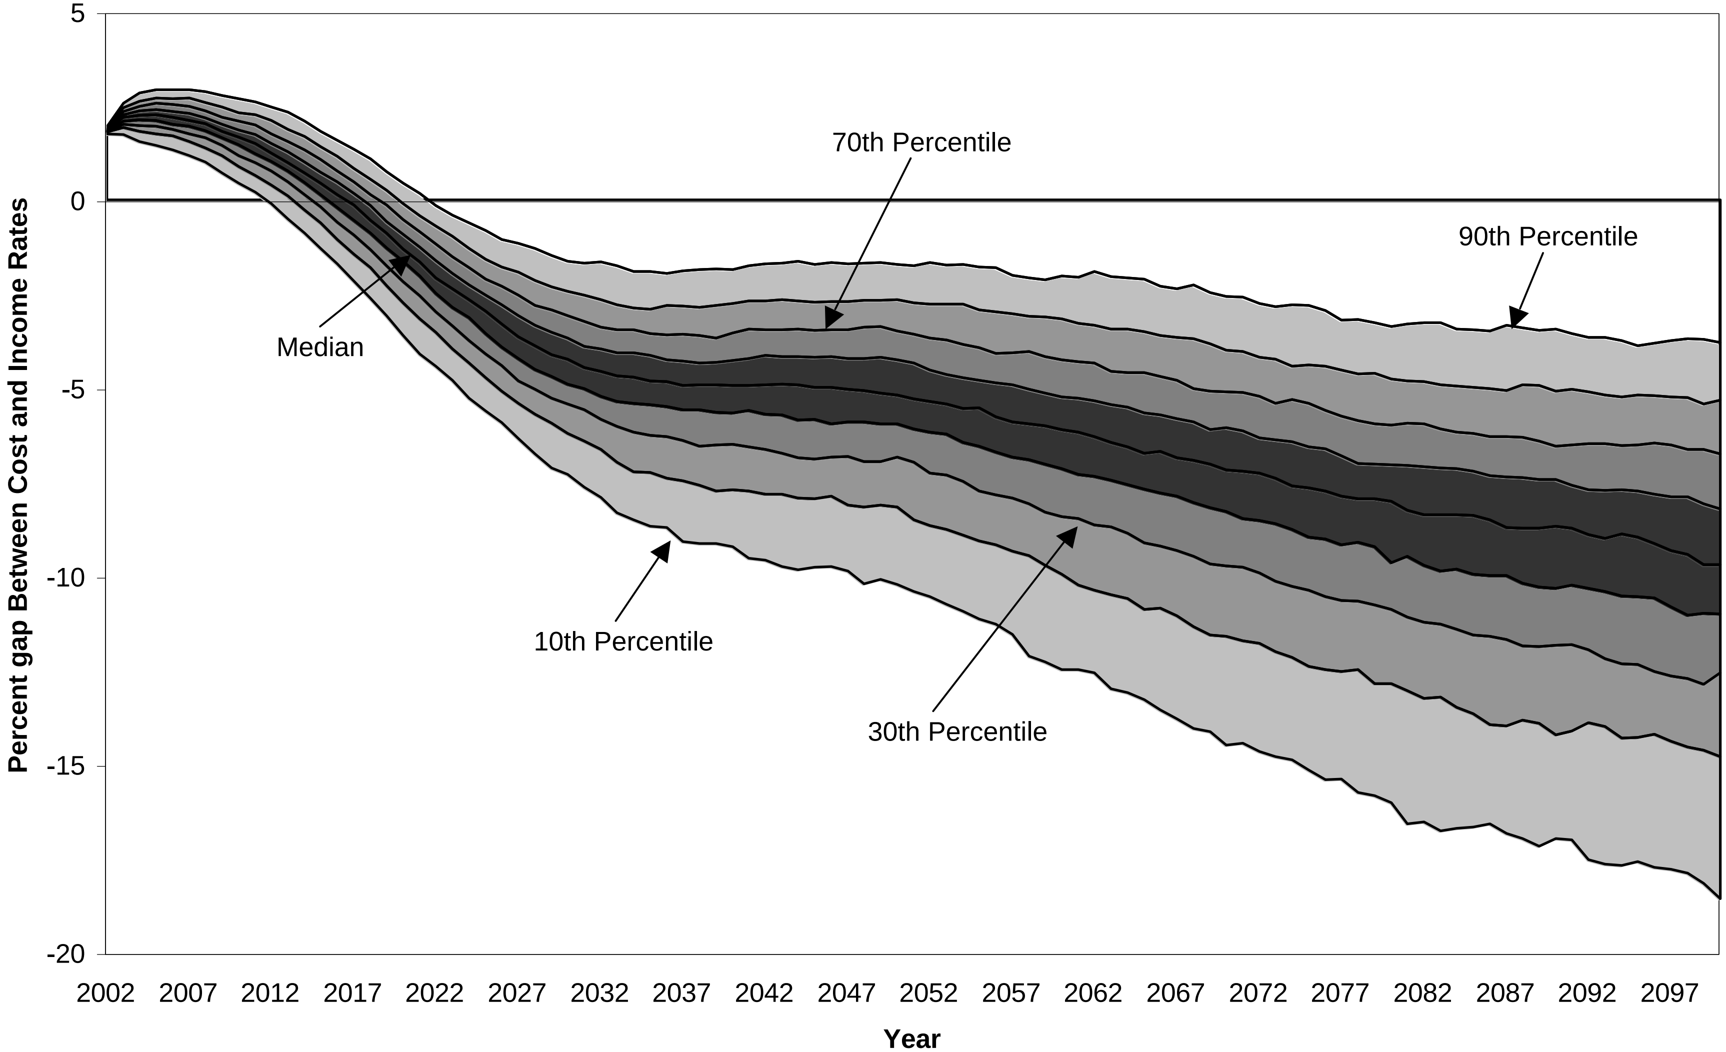

Figure 4: Distribution of Trust Fund Ratios
Current Law, No Fail-Safe (Insolvent System)

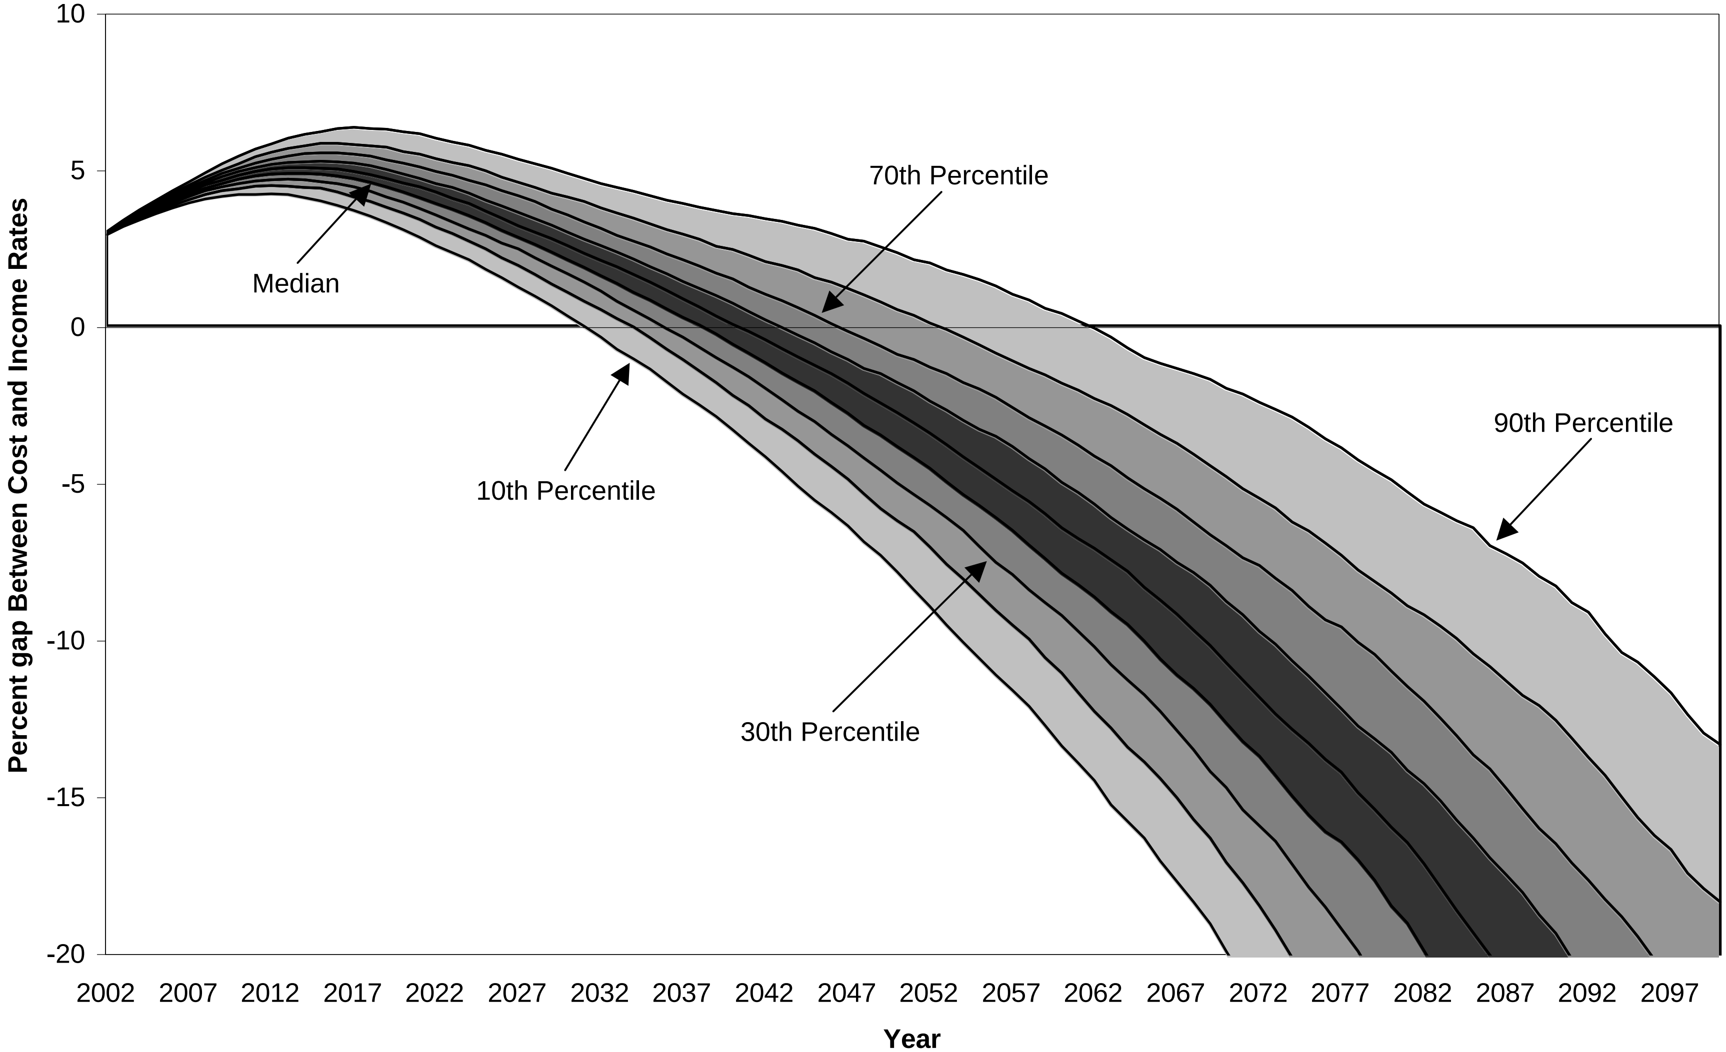

Figure 5: Median Internal Rates of Return by Birth Cohort (Average Workers Under Current Law)

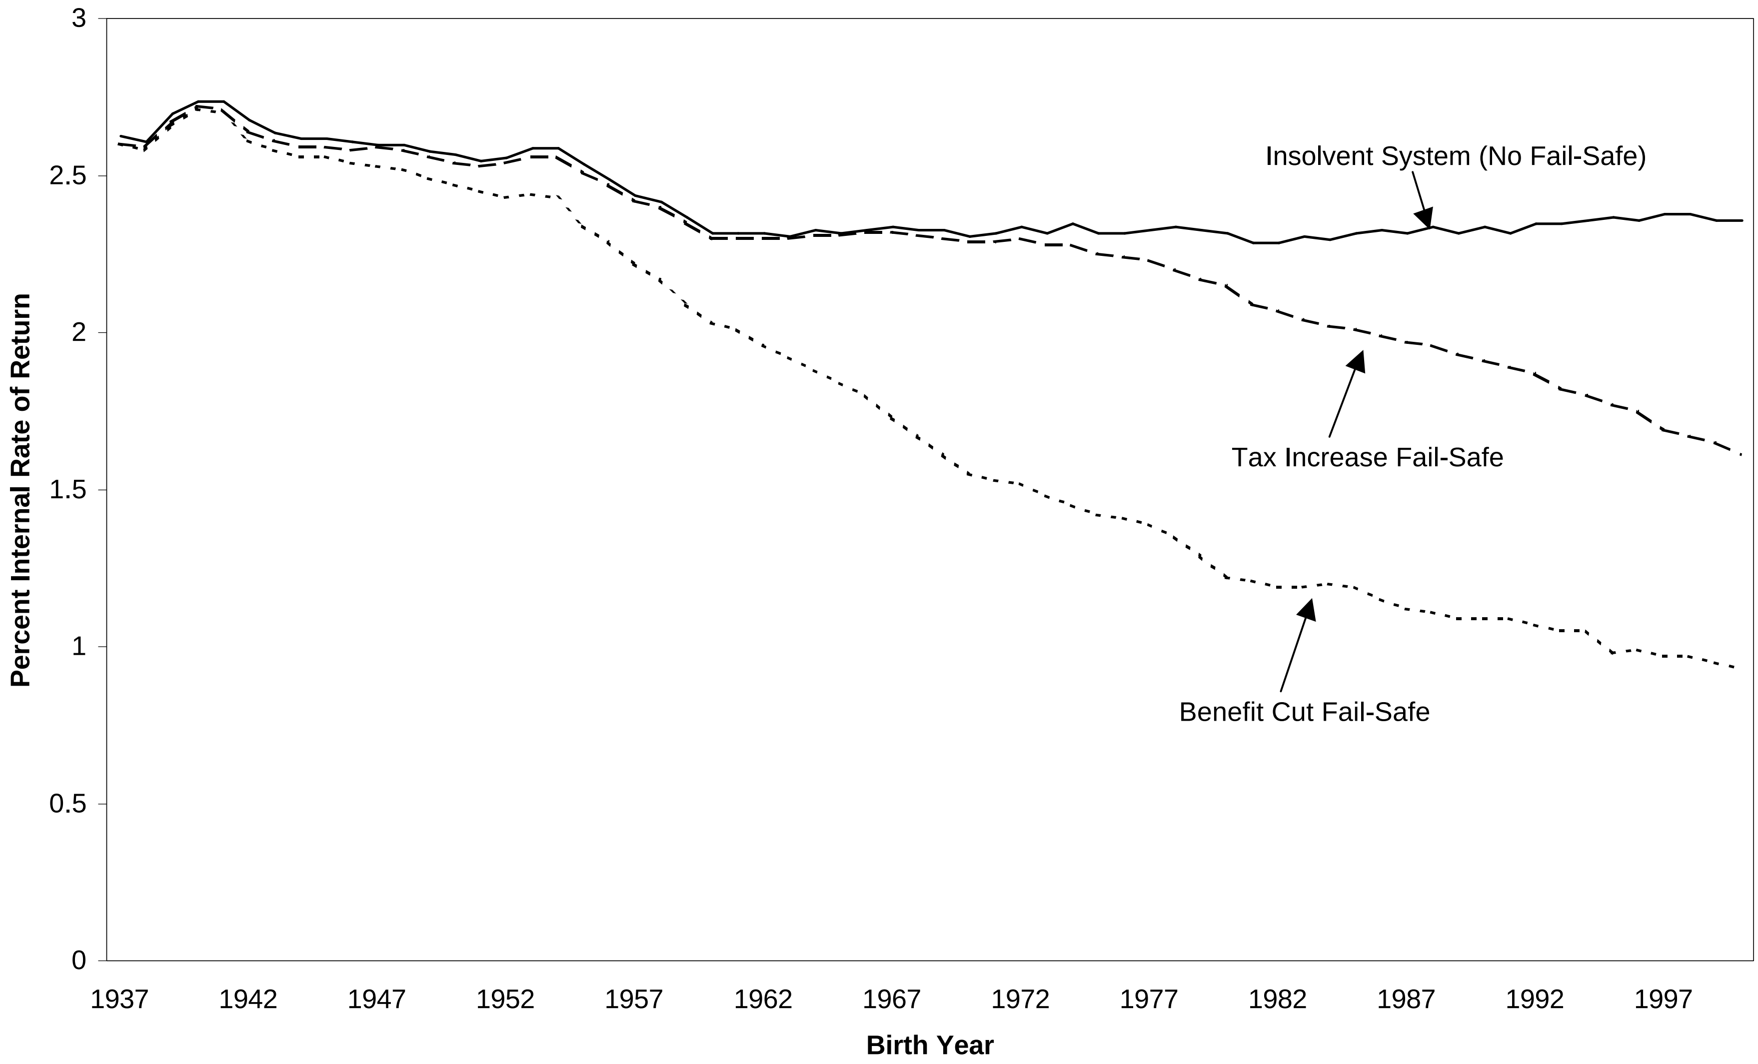

Figure 6: Probability Distribution of Internal Rates of Return for Average Workers
Current Law with Benefit Cut Fail-Safe

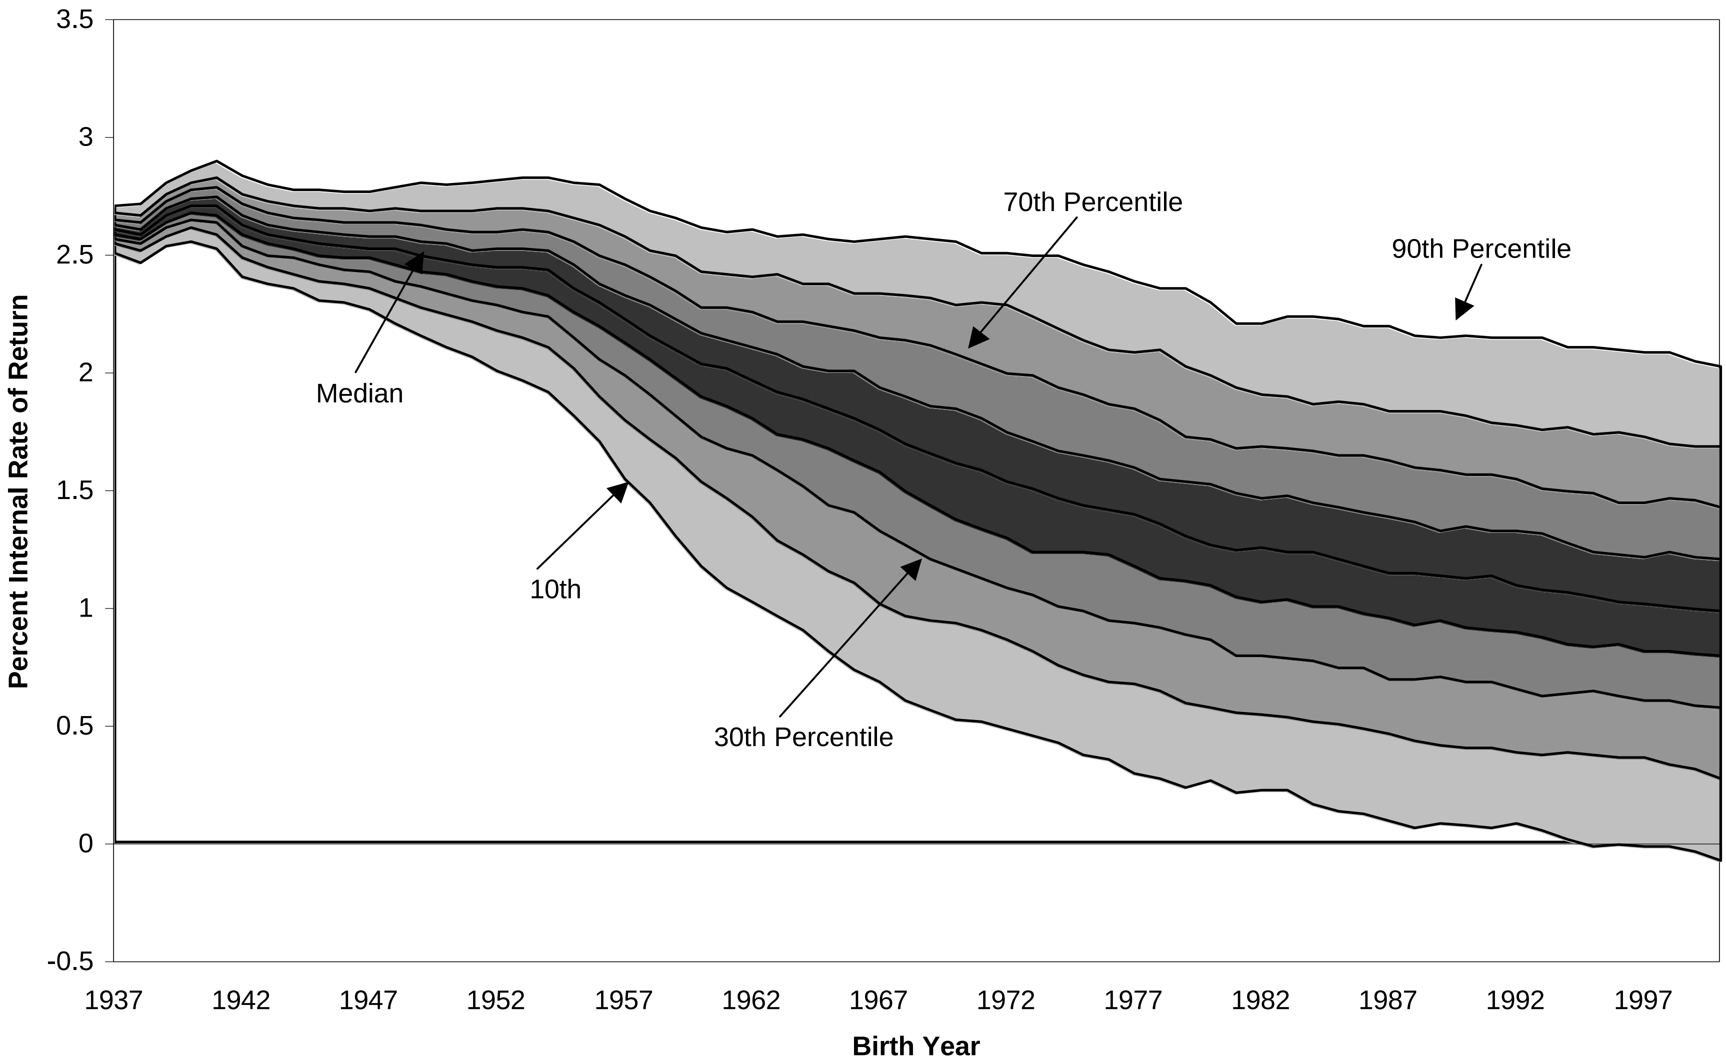

Figure 7: Probability Distribution of Benefit Cuts
Current Law with Benefit Cut Fail-Safe

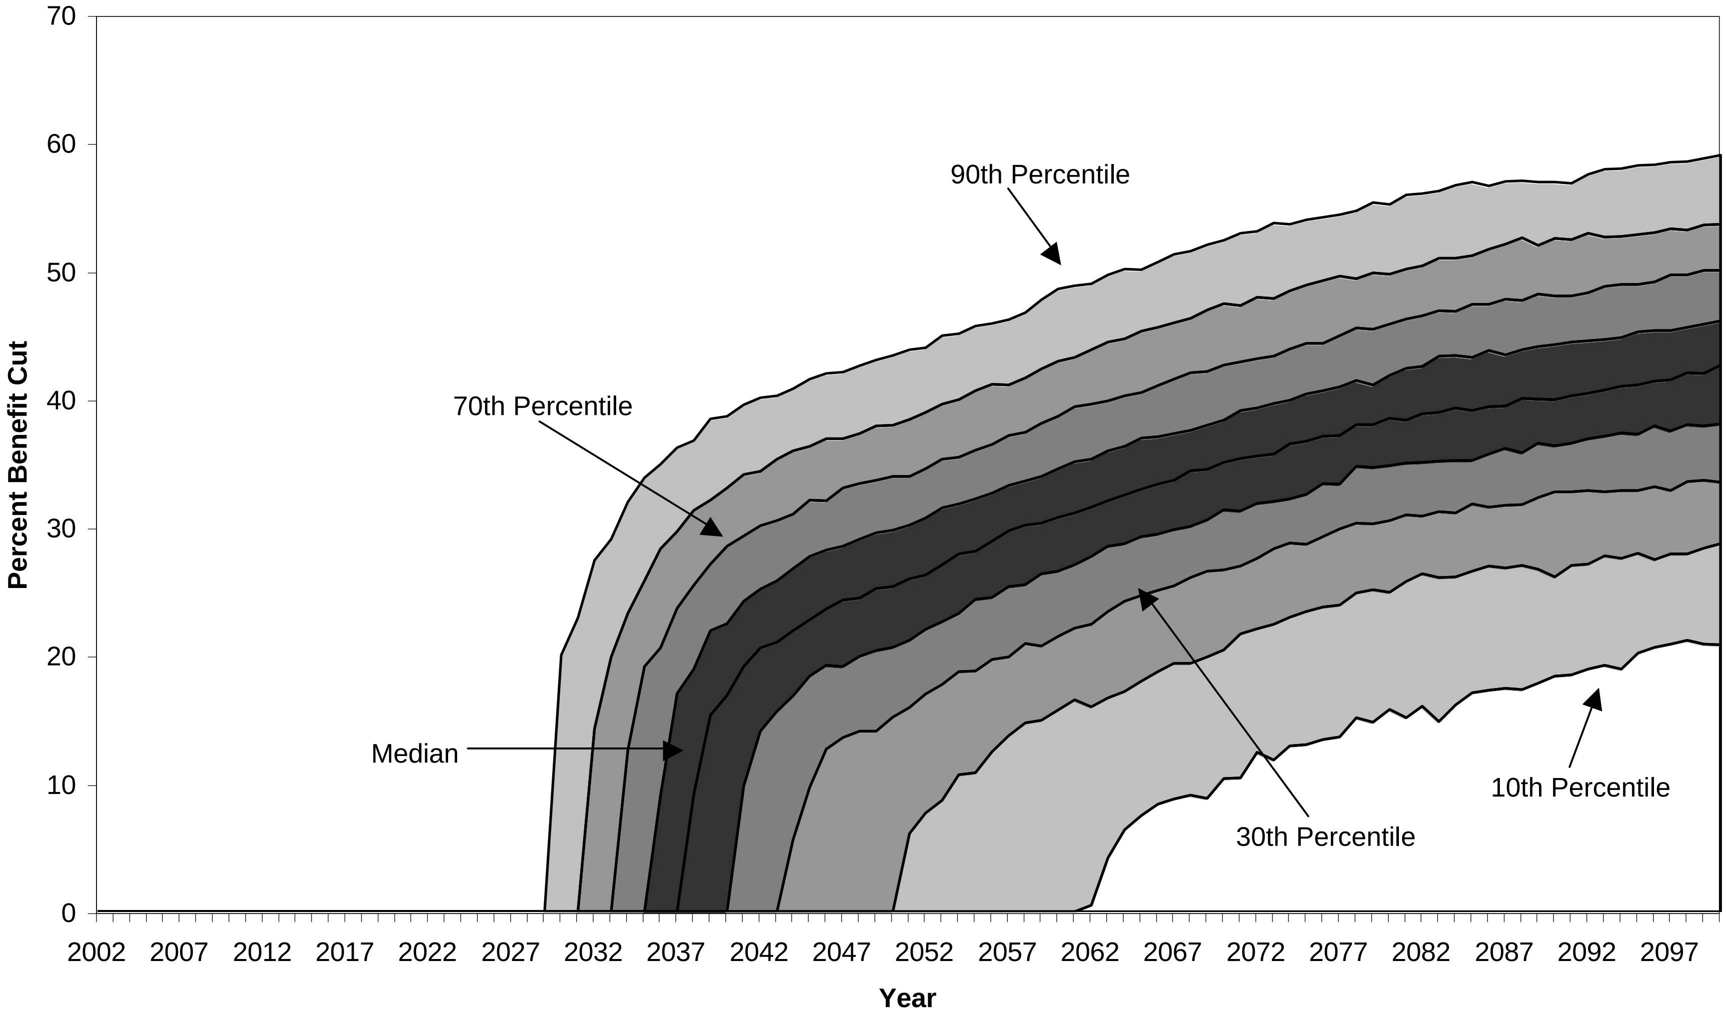

Figure 8: Probability Distribution of Internal Rates of Return for Average Workers Current Law with Tax Increase Fail-Safe

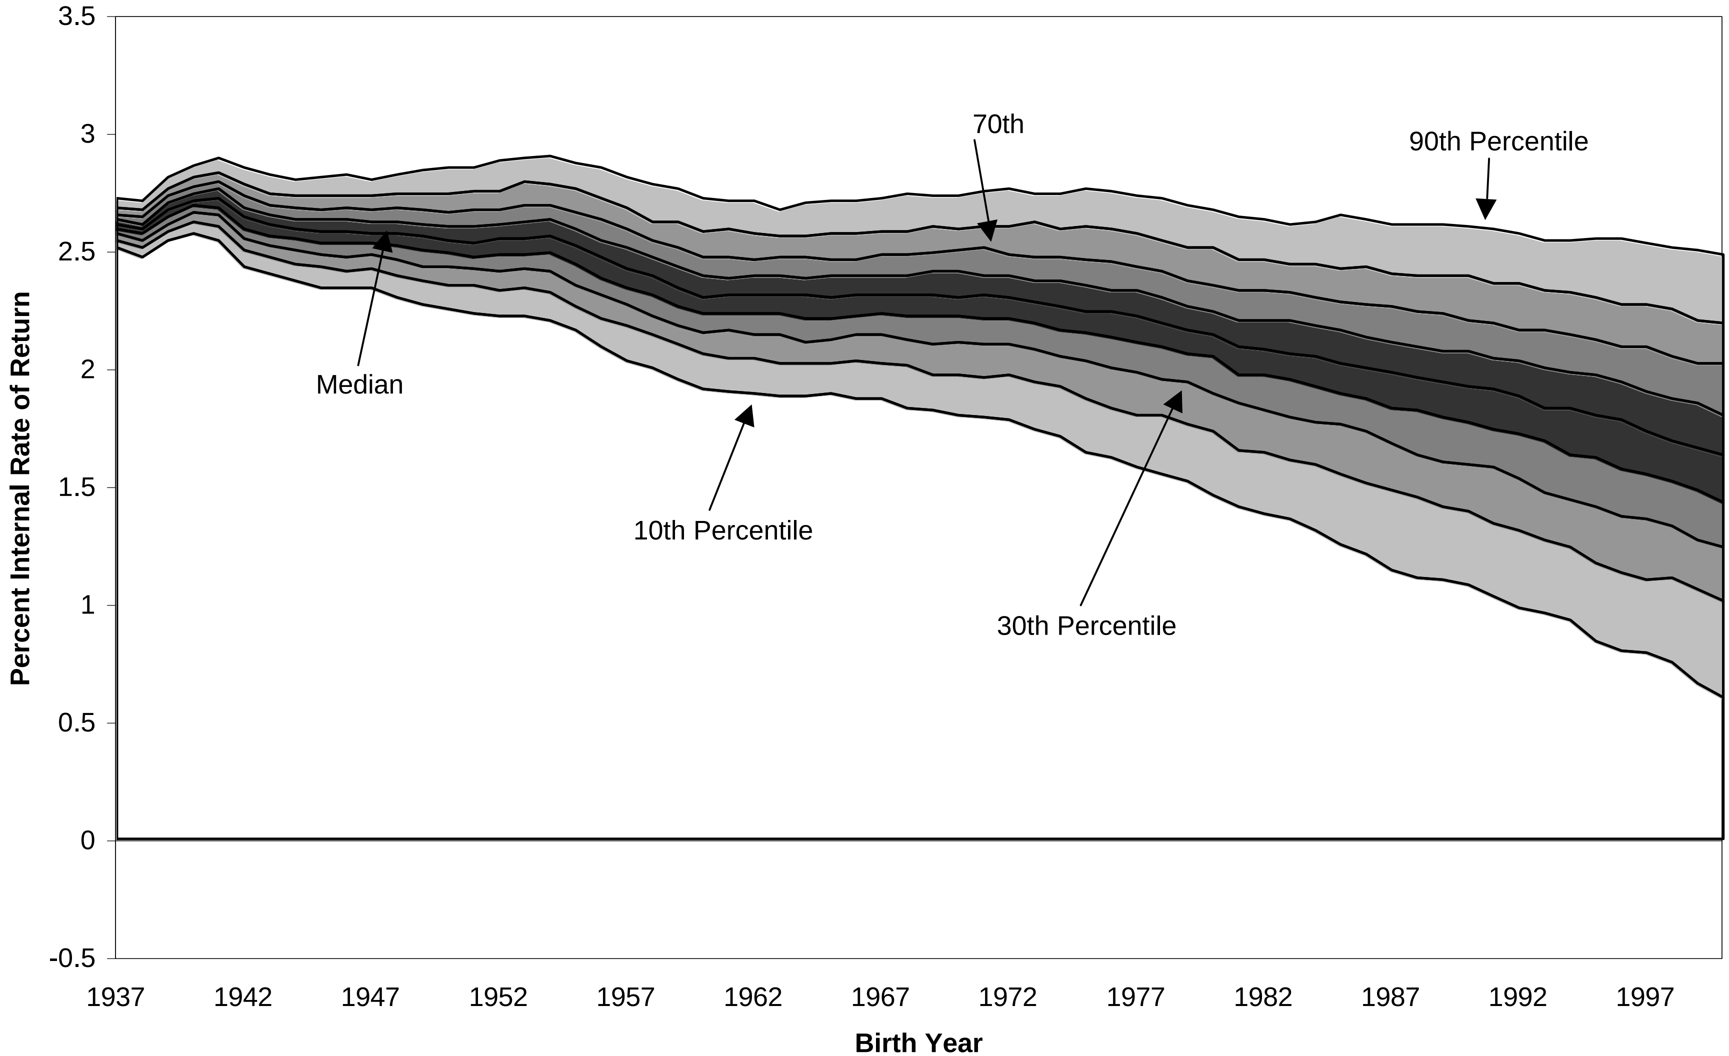

Figure 9: Probability Distribution of Tax Increases
Current Law with Tax Increase Fail-Safe

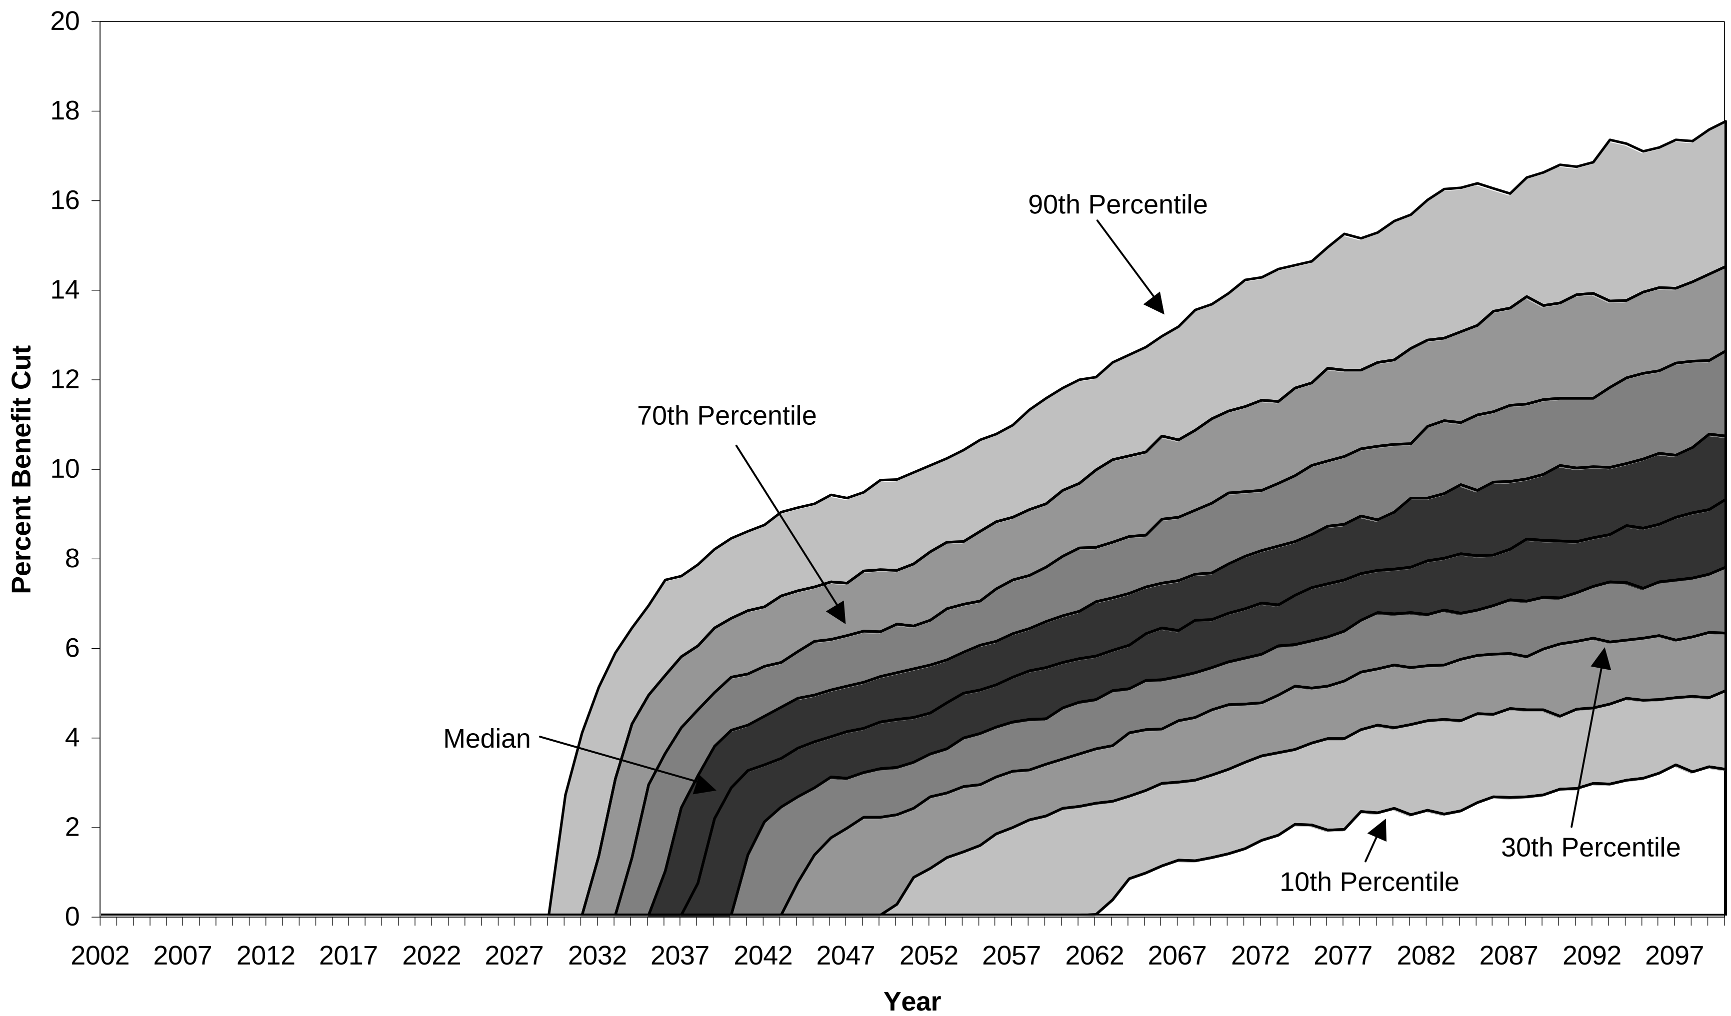

Figure 10: Life Expectancy at Age 65 by Calendar Year

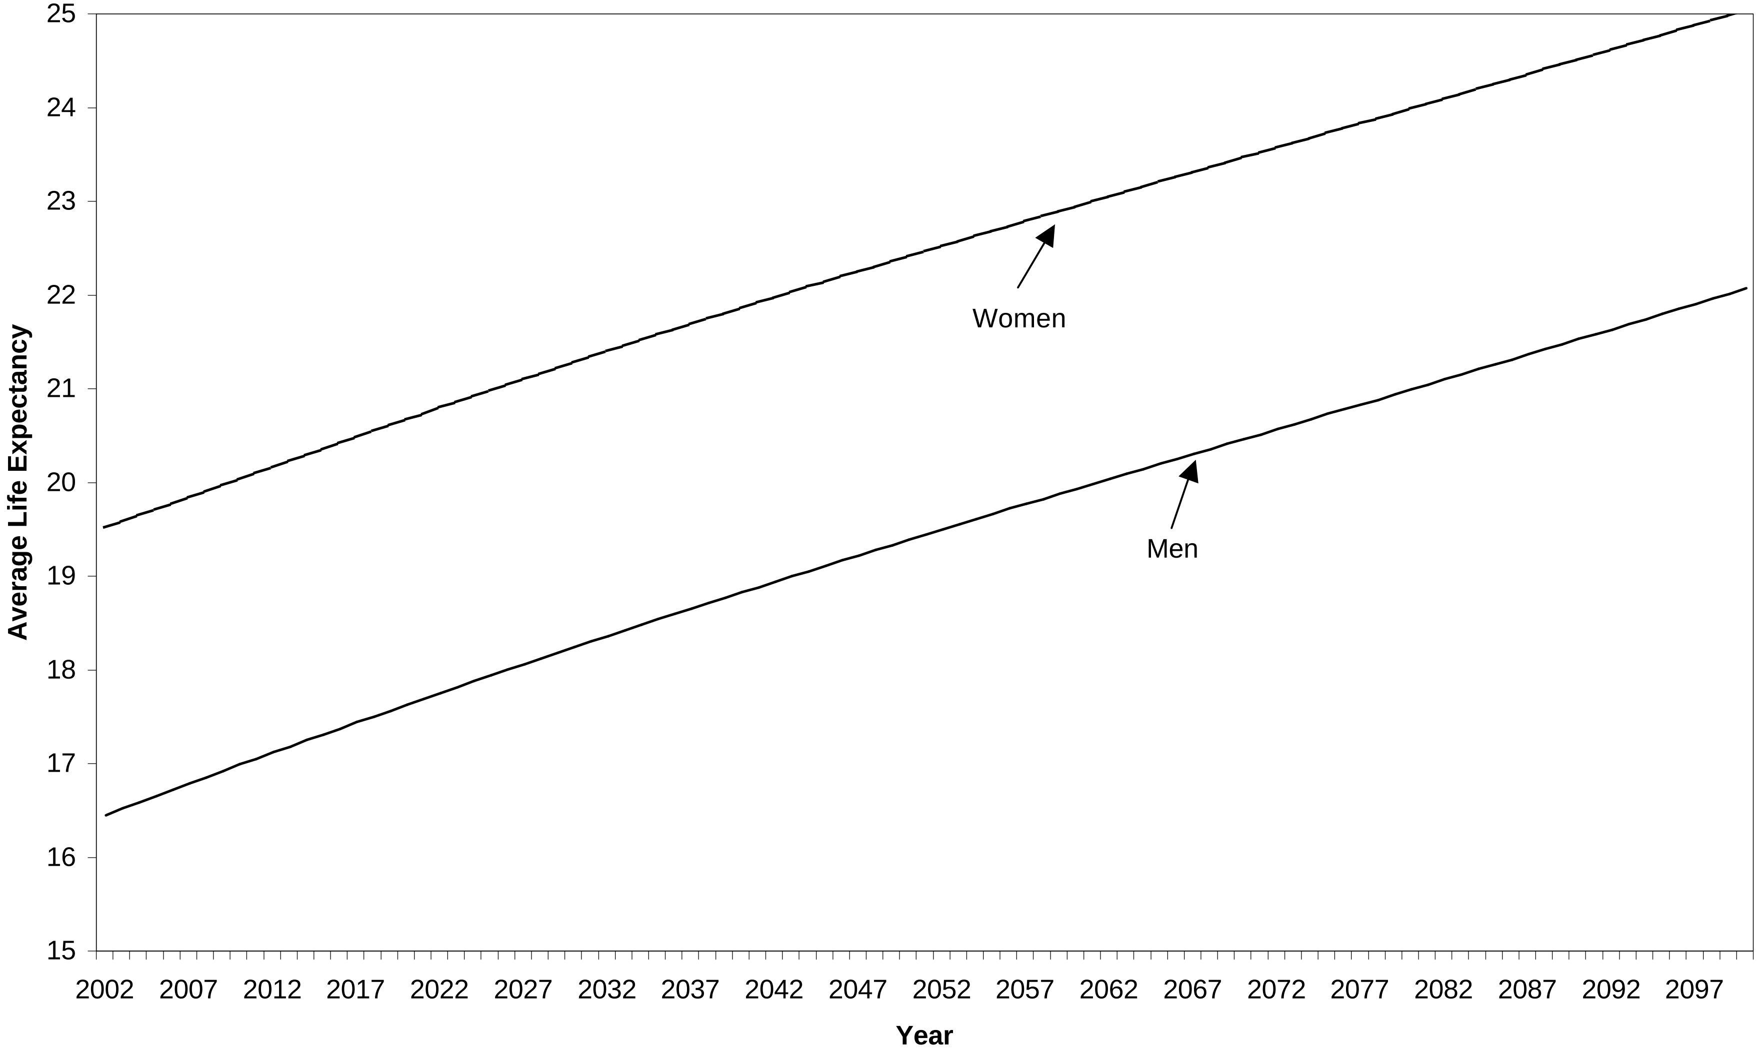

Figure 11: Distribution of Gap Between Cost and Income Rates
Normal Retirement Age Increase to 70, No Fail-Safe (Insolvent System)

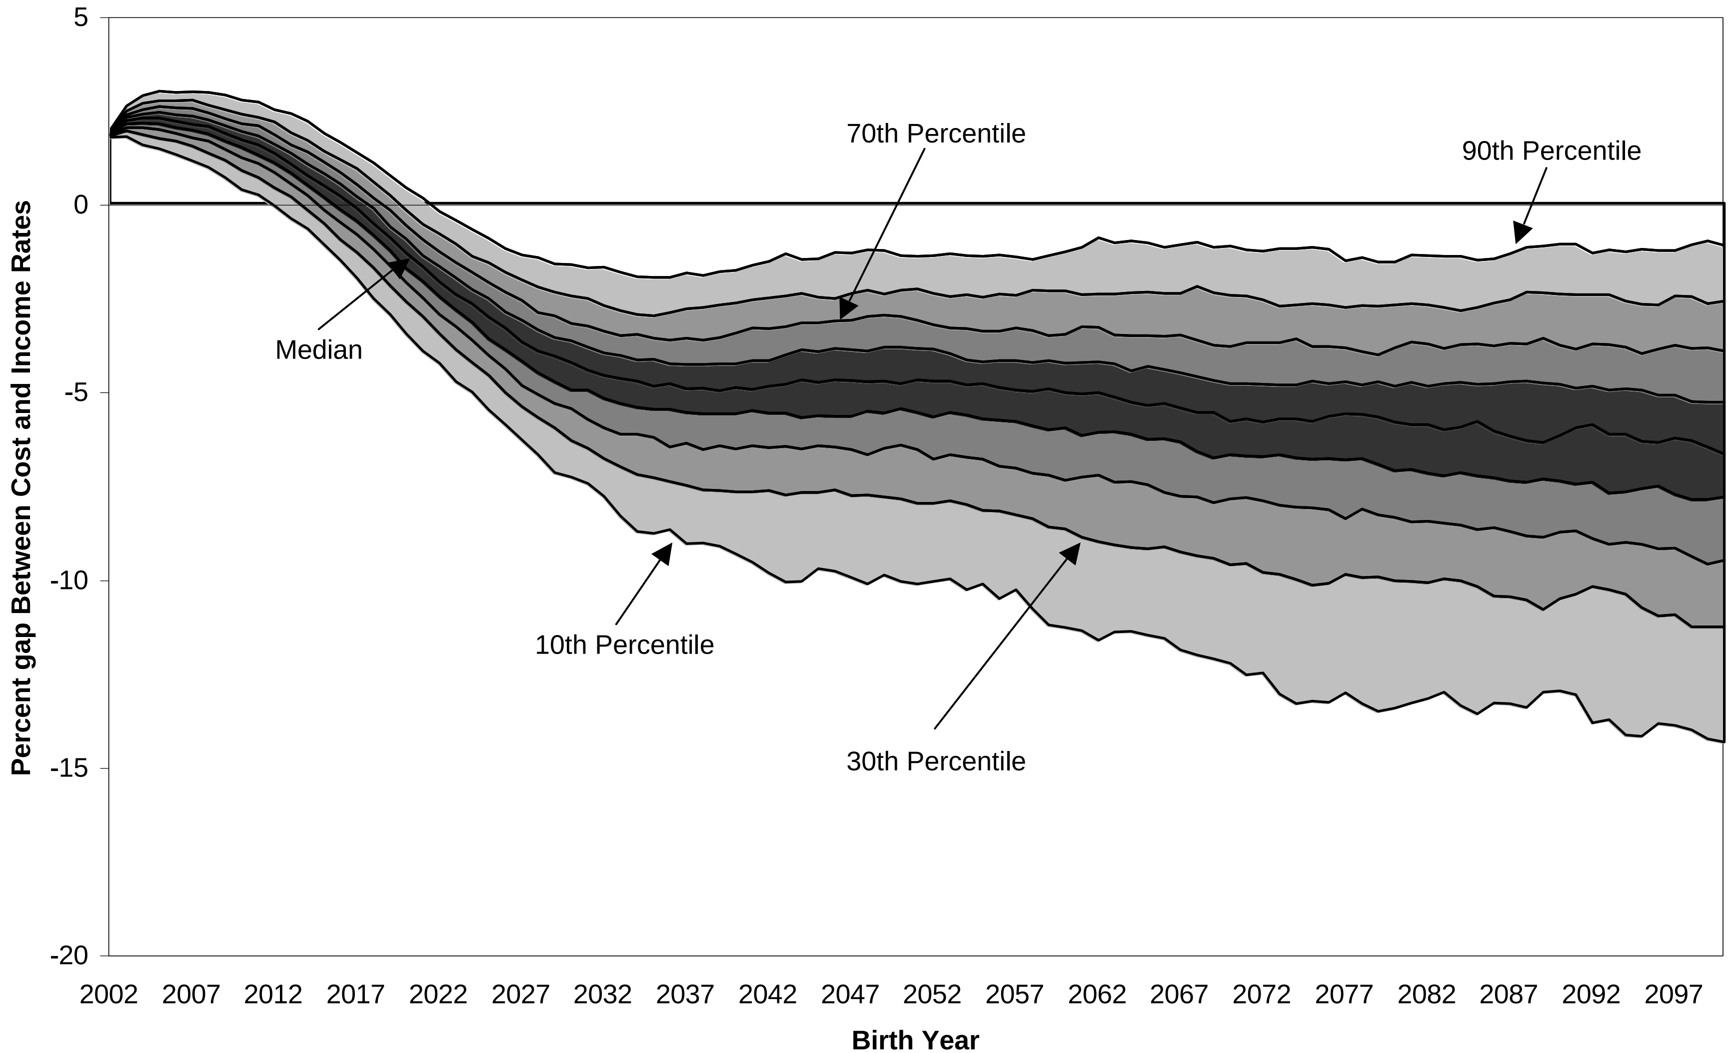

Figure 12: Distribution of Trust Fund Ratios
Normal Retirement Age Increase to 70, No Fail-Safe (Insolvent System)

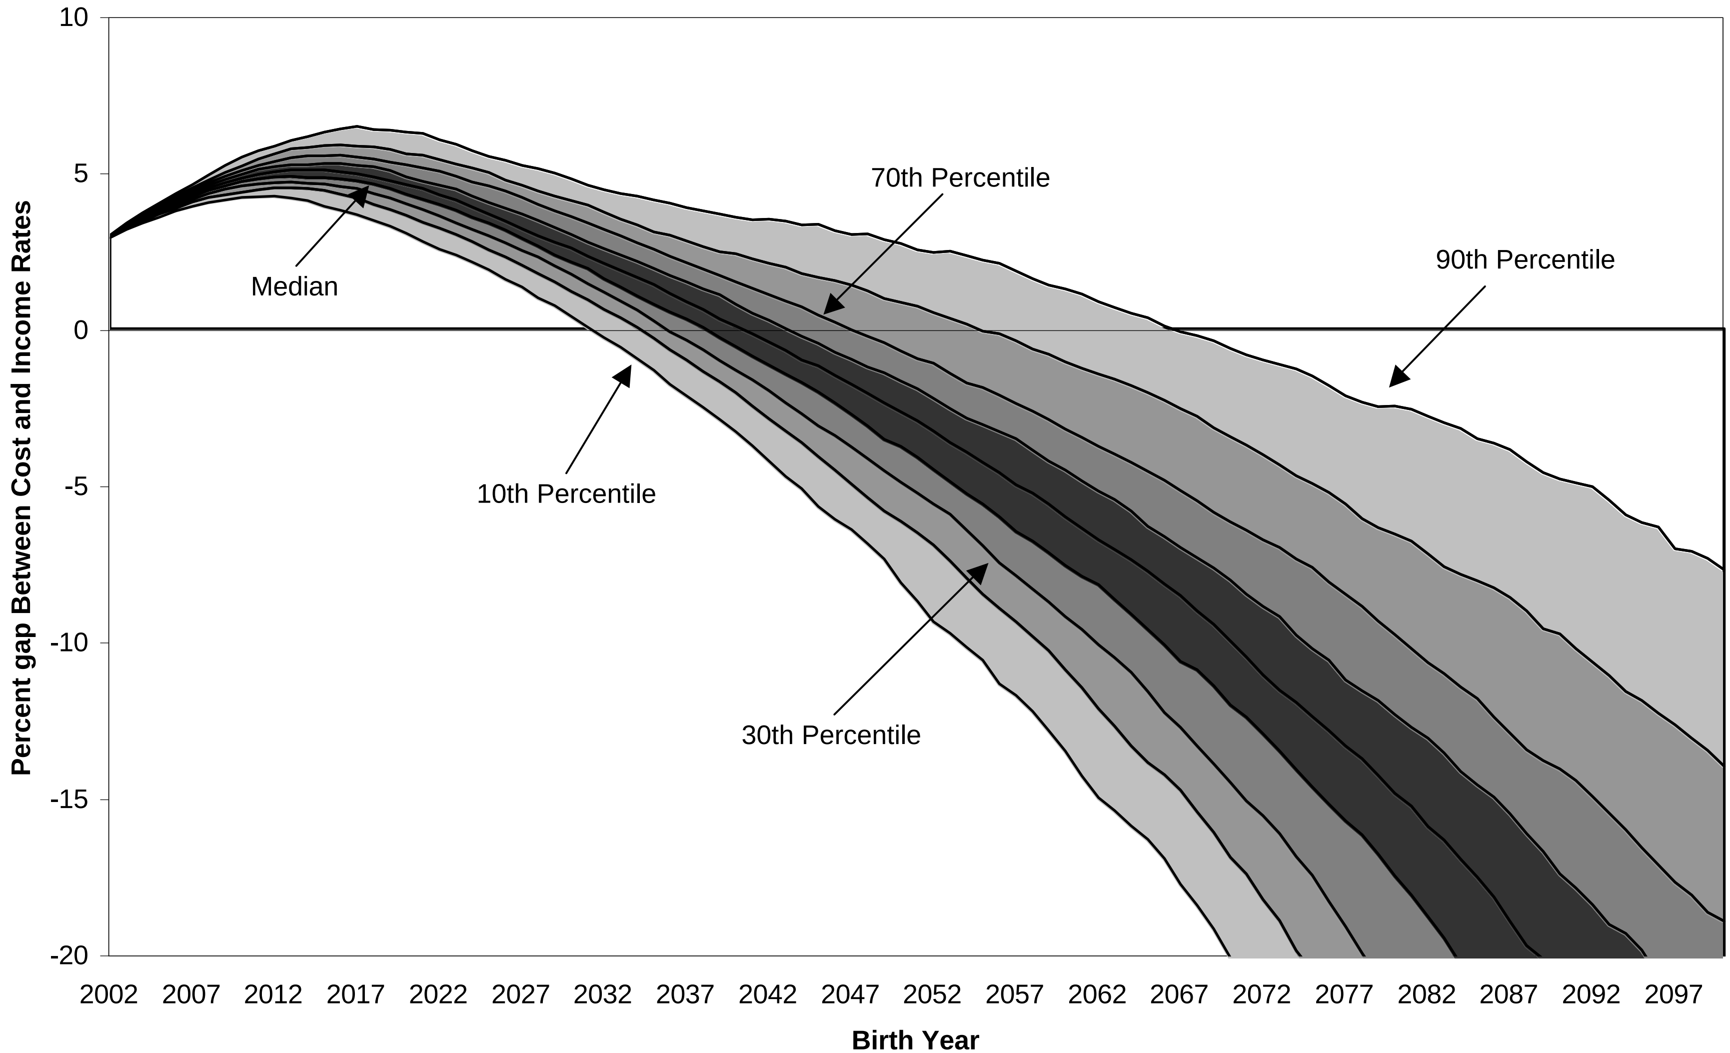

Figure 13: Median Internal Rates of Return by Birth Cohort
(Average Workers Under Current Law and Increased Normal Retirement Age to 70)

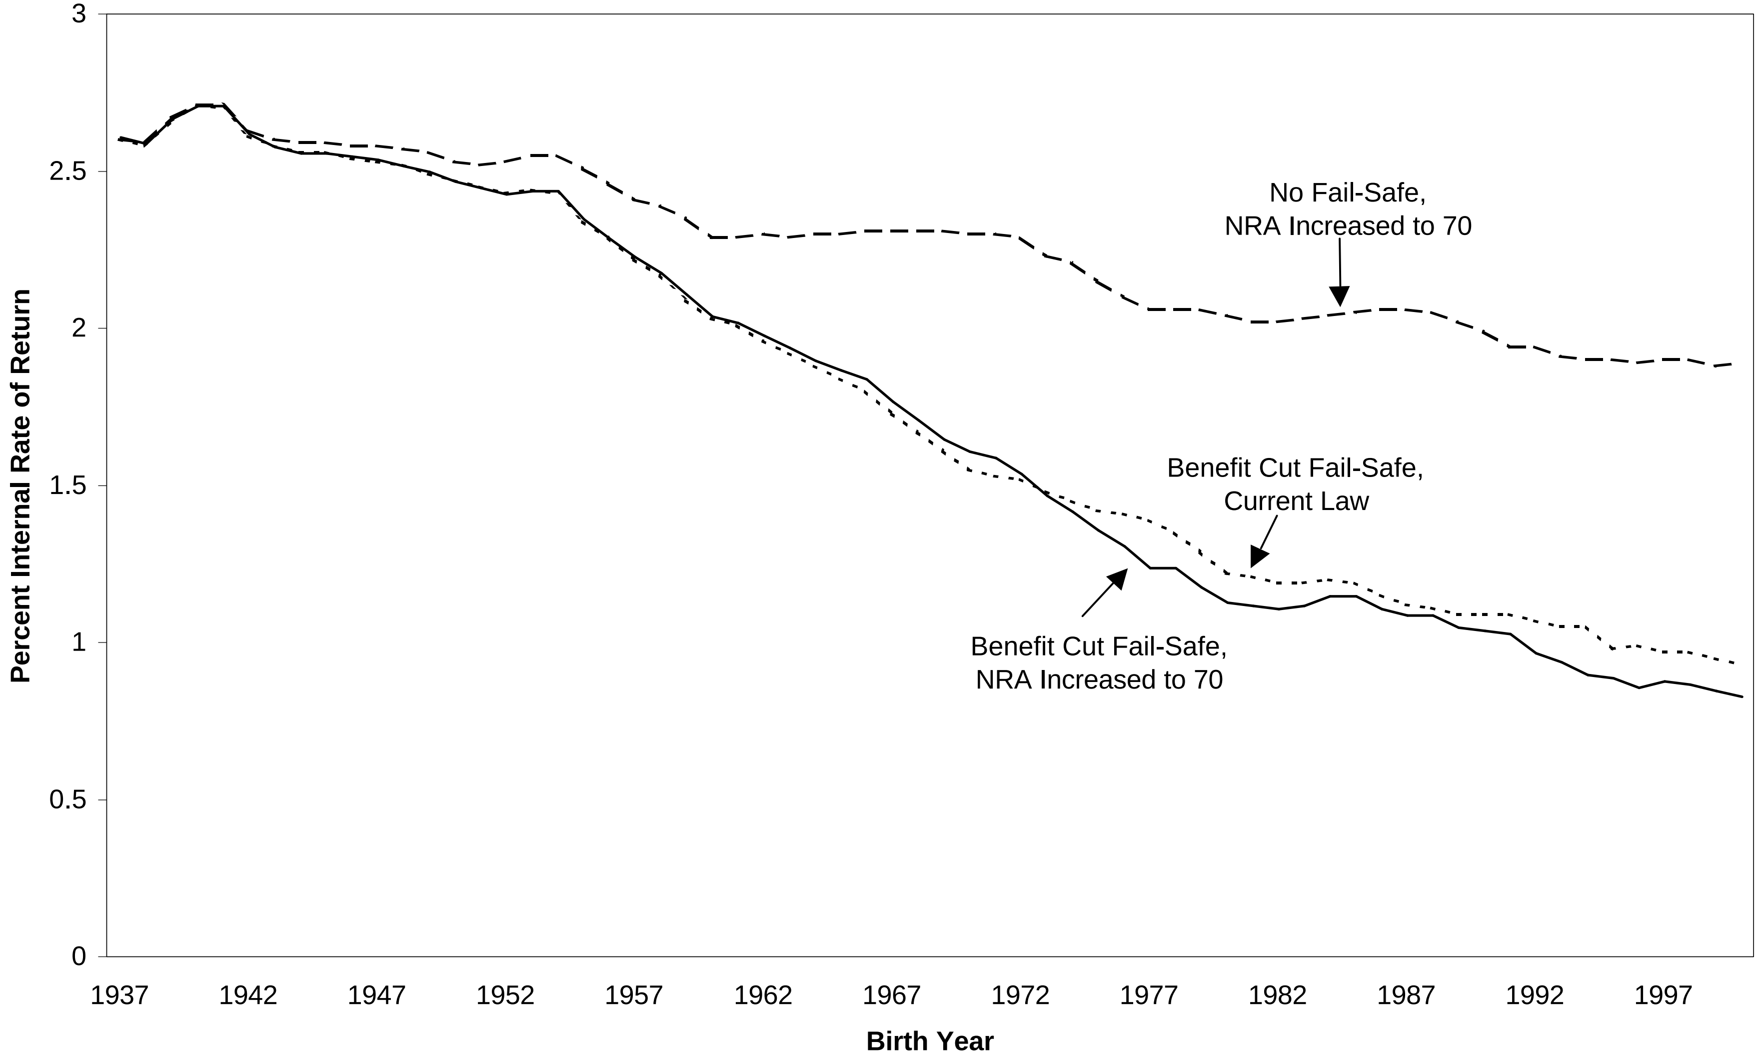

Figure 14: Median Internal Rates of Return by Birth Cohort
(Average Workers, Increased NRA and Benefit Cut Fail-Safe, Various Reforms)

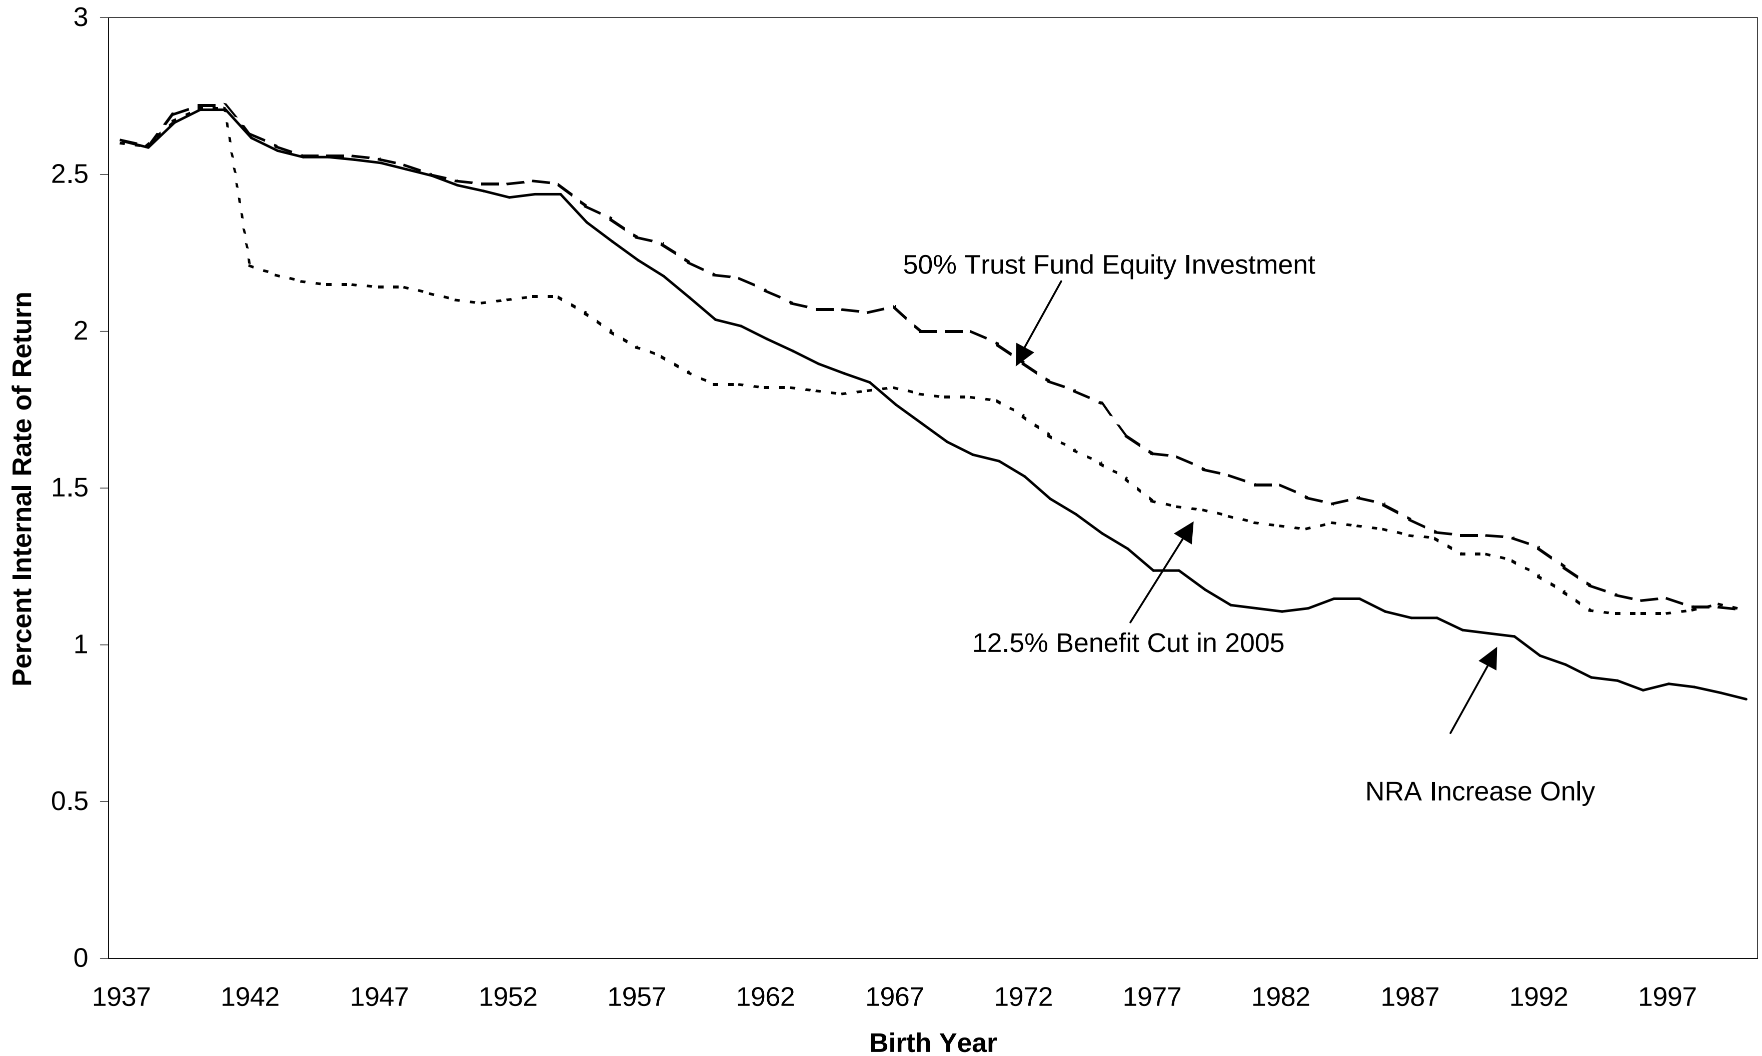

Figure 15: Median Internal Rates of Return by Birth Cohort
(Average Workers, Increased NRA and Benefit Cut Fail-Safe, Various Reforms)

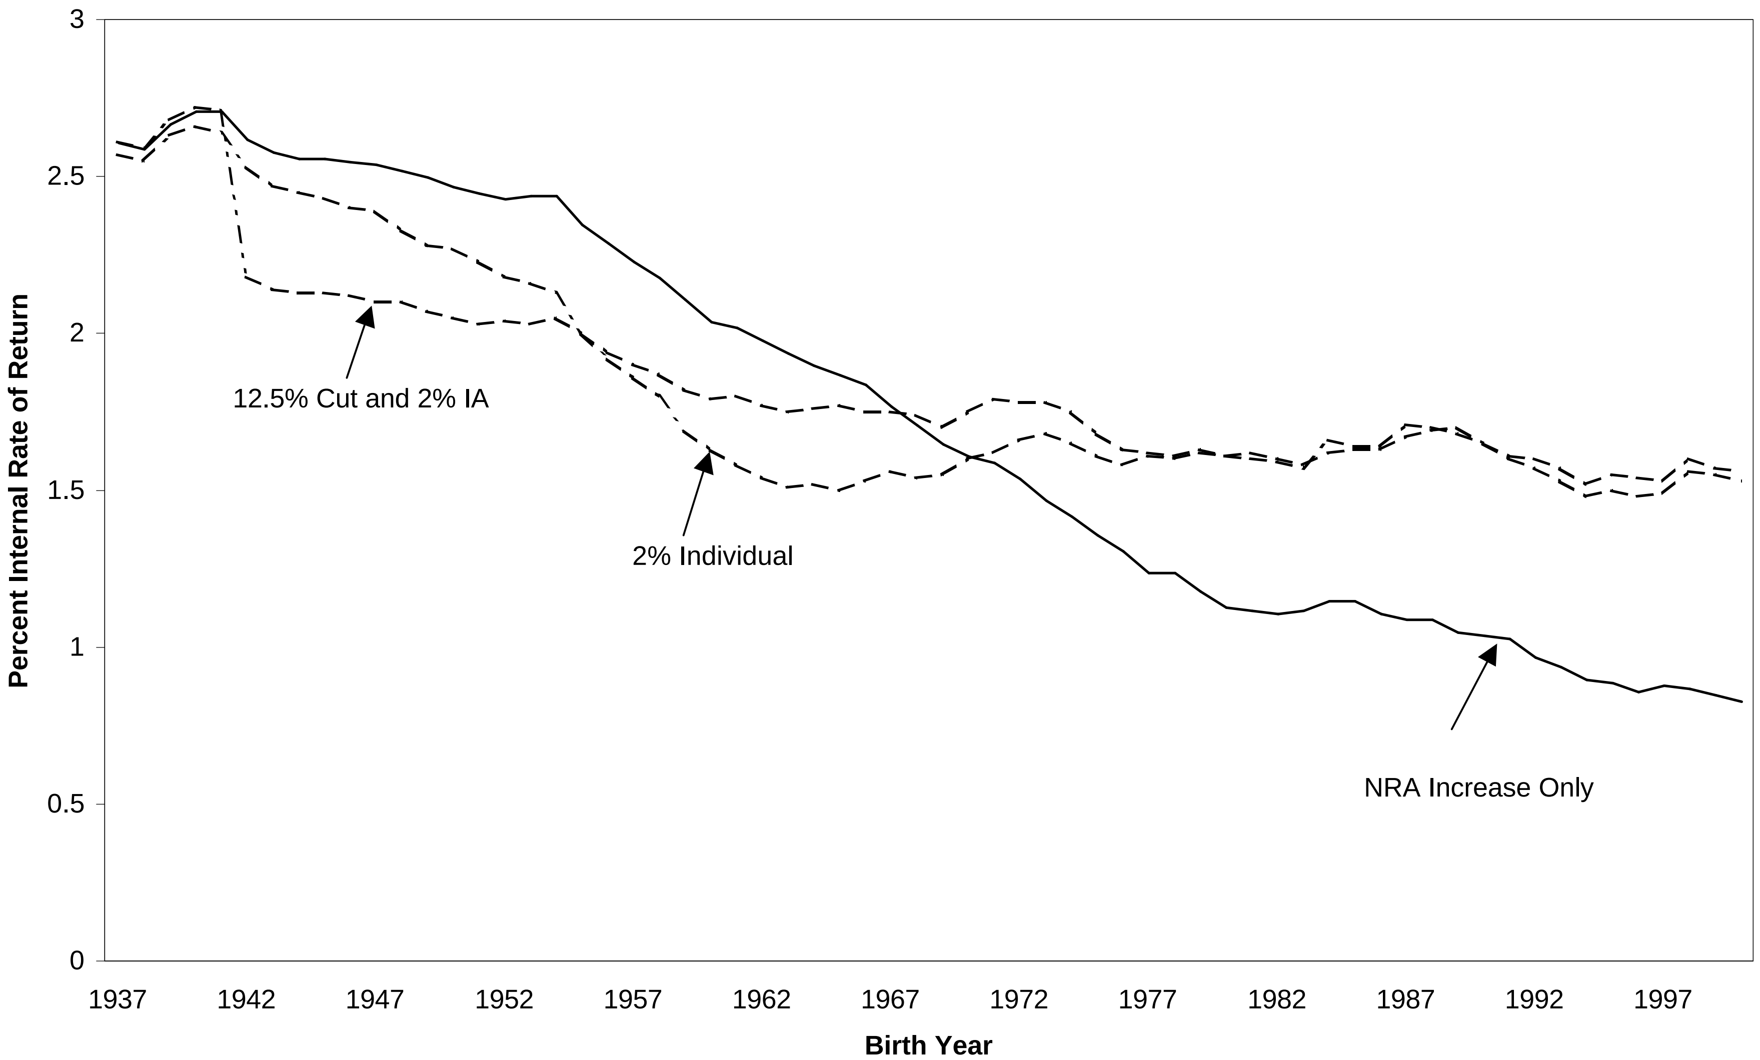

Table 3
IRR Deciles for Average Earners by Birth Cohort Under Various Reform Scenarios
(All Simulations Start with NRA Increase to Age 70 and Benefit Cut Fail-Safe)

	Percentile of IRR Probability Distribution				
	10th	**30th**	**Median**	**70th**	**90th**
1940 Birth Cohort					
Base Reform (NRA Change Only)	2.54	2.63	2.70	2.76	2.86
12.5% New Award Reduction Starting 2007	2.56	2.65	2.71	2.77	2.85
50% Trust Fund Equity Investment	2.56	2.66	2.72	2.77	2.86
2% Individual Account Carve Out	2.45	2.58	2.66	2.73	2.80
12.5% Award Reduction and 2% Individual Account	2.54	2.65	2.72	2.76	2.85
1955 Birth Cohort					
Base Reform (NRA Change Only)	1.78	2.11	2.34	2.54	2.79
12.5% New Award Reduction Starting 2007	1.64	1.88	2.06	2.22	2.42
50% Trust Fund Equity Investment	1.83	2.22	2.40	2.60	2.81
2% Individual Account Carve Out	1.24	1.69	2.00	2.31	2.72
12.5% Award Reduction and 2% Individual Account	1.33	1.75	2.00	2.22	2.46
1970 Birth Cohort					
Base Reform (NRA Change Only)	0.55	1.19	1.60	2.10	2.54
12.5% New Award Reduction Starting 2007	0.83	1.48	1.79	2.02	2.28
50% Trust Fund Equity Investment	0.75	1.47	2.00	2.31	2.58
2% Individual Account Carve Out	0.50	1.22	1.60	2.05	2.55
12.5% Award Reduction and 2% Individual Account	0.58	1.30	1.75	2.13	2.53
1985 Birth Cohort					
Base Reform (NRA Change Only)	-0.03	0.66	1.14	1.58	2.06
12.5% New Award Reduction Starting 2007	0.09	0.89	1.38	1.66	2.00
50% Trust Fund Equity Investment	0.20	0.90	1.47	1.88	2.28
2% Individual Account Carve Out	0.37	1.07	1.56	1.97	2.55
12.5% Award Reduction and 2% Individual Account	0.41	1.18	1.63	2.04	2.62
2000 Birth Cohort					
Base Reform (NRA Change Only)	-0.32	0.38	0.82	1.28	1.80
12.5% New Award Reduction Starting 2007	-0.25	0.53	1.11	1.45	1.83
50% Trust Fund Equity Investment	-0.17	0.57	1.11	1.65	2.08
2% Individual Account Carve Out	0.13	1.07	1.56	1.97	2.55
12.5% Award Reduction and 2% Individual Account	0.13	1.02	1.53	1.89	2.50