Working Paper Series
Congressional Budget Office
Washington, D.C.

UNCERTAIN POLICY
FOR AN UNCERTAIN WORLD:
THE CASE OF SOCIAL SECURITY

John Sabelhaus (e-mail: john.sabelhaus@cbo.gov)
Julie Topoleski (e-mail: julie.topoleski@cbo.gov)
Congressional Budget Office
Washington, D.C.

March 2006
2006-05

Abstract

Analysis and discussion of Social Security policy are usually based on expected fiscal and societal outcomes. However, future demographic and economic trends are uncertain, and thus ultimate outcomes for aggregate system financial flows and the distribution of taxes and benefits across generations are uncertain. This paper analyzes a state-dependent approach to policy in which future Social Security benefit formulas are tied to realized economic and demographic outcomes over time. The results, based on a microsimulation model with stochastic capabilities, show the extent to which it is possible to systematically address uncertainty about system finances and distributional outcomes.

1. Introduction[1]

Although it is widely acknowledged that the U.S. Social Security system is likely to become insolvent within a few decades, the eventual realized imbalance between inflows and outflows will ultimately depend on demographic and economic factors whose values are not yet known.[2] When confronting this type of uncertainty, one approach to policy is to establish rules for modifying benefit formulas when demographic and economic outcomes are realized, rather than to try to choose fixed-benefit formulas that are expected to hold for the indefinite future.[3] An example, and one of the policy options explored in this paper, is to systematically lower benefit levels (at a given retirement age) as longevity increases over time. The goal of this paper is to quantify the effects of tying benefit changes to realized economic and demographic trends, where the outcomes of interest are the time path of Social Security system finances and the distribution of taxes and benefits across birth cohorts.

Consider the effect of a schedule of benefit cuts chosen to achieve expected solvency for the indefinite future: The expected cost of achieving system solvency would be distributed across generations depending on the timing of the benefit cuts, and normative judgments would play a role in deciding how that cost should be allocated. However, system solvency and distributional outcomes are uncertain. Indeed, if the models and assumptions are correct, there is

---

[1]An earlier version of this paper was presented at the Association of Public Policy Analysis and Management Annual Conference in Washington, D.C. (November 2005). The authors would like to thank Julian Cristia, Amy Rehder Harris, Noah Meyerson, Michael Simpson, and Reuben Snipper for useful comments and suggestions.

[2]See Social Security Administration Board of Trustees (2005) and Congressional Budget Office (2004).

[3]This argument is consistent with the more general analysis of policymaking under uncertainty discussed by Popper, Lempert, and Bankes (2005), and the specific example of state budgeting under revenue uncertainty analyzed by Cornia, Nelson, and Wilko (2004).

a 50 percent chance that the benefit cuts chosen to achieve solvency would be too large; there is also a 50 percent chance benefits would be cut too little, which would result in shortfalls. Those cuts would be larger than needed in good budgetary states of the world, because those good states of the world would be reflected in a Social Security surplus greater than expected. Although this inherent uncertainty cannot be eliminated, it can be quantified, and can be considered when balancing policy objectives.

The two sets of policy objectives analyzed here are aggregate system finances and distributional outcomes across generations. Many would agree that the Social Security system should be put on a sustainably solvent trajectory, one in which benefits will continue to be paid without interruption from financing shortfalls. There are many Social Security policy options that meet the sustainable solvency criteria, but they have very different distributional effects. For example, one could impose significant immediate benefit cuts and build up a Social Security trust fund that is so large the interest on the fund would allow (the now lower) benefits to be paid indefinitely.[4] At the other extreme, one could schedule future benefit cuts to occur only as needed to keep benefits payable on a flow basis. Thus, the trust fund would fall to zero, there would be no interest earnings, and future benefits paid would be limited by taxes collected. This is functionally equivalent to the current-law policy of allowing trust fund depletion to trigger the timing of benefit cuts, but, in principle, one could legislate exactly that pattern of cuts, and thus it serves as a useful benchmark.

These extreme approaches to putting the system on a solvent trajectory have very

---

[4]This abstracts from any interaction between on-budget and off-budget deficits. For example, Smetters (2004) argues that it cannot be assumed that trust fund surpluses will be saved for future benefit payments, and so pre-funding the Social Security trust funds does not necessarily reduce the level of debt held by the public.

different distributional implications across birth cohorts because of the timing of benefit changes. There is of course a middle ground, which is to initiate benefit cuts before the trust fund is exhausted, but not to cut benefits for near-term retirees so much that they pay a disproportionate share of the costs. The key insight of this paper is that the middle ground approach can be taken one step further: Tying the extent and timing of benefit changes to realized demographic and economic outcomes may be a way to address the competing objectives of distributional fairness and maintenance of a specific funding strategy.

Two state-dependent policies are analyzed in which future benefits depend on realized longevity and productivity trends over time. The first policy change ties future benefit cuts to longevity increases. As life spans increase for future cohorts, benefits (at any given benefit-claiming age) would be systematically reduced so that lifetime benefits would not increase relative to lifetime taxes paid. The second policy change modifies growth factors applied to benefits after benefit claiming begins. Rather than just growing benefits with inflation as in the current system, benefits would also be tied to productivity, rising faster than inflation when productivity is above its expected value, and vice versa. Both the longevity-adjusted benefits and productivity-adjusted cost of living policies shift uncertainty from future taxpayers onto future beneficiaries.

This paper quantifies the effects of these two state-dependent policies on system finances and distributional outcomes. The two policies are compared to a base scenario in which benefits are cut only once by the amount needed to achieve expected sustainable solvency.[5] The state-

[5]The illustrative expected-solvency scenario involves an across-the-board 19 percent cut in new benefit awards starting in 2012. The year of the benefit cut (which is the same for the two policy experiments) is chosen so that no one older than age 55 in 2005 would be affected, which reflects recent reform proposals that have been put forth. The other key assumption of the scenario is that the benefit cut is not reversible; that eliminates situations

dependent policies also achieve expected solvency, but the initial benefit cut required is smaller than in the base case because of subsequent longevity indexing. Because the initial benefit cut is smaller, and because further cuts are tied to (stochastic) realized economic and demographic outcomes, the state-dependent policies effectively shift a significant fraction of the upside potential for future trust funds into upside potential for beneficiaries. Benefit cuts are not invoked in states of the world when they are not needed, for example, when mortality improvement is low and productivity growth is high. Further, although there is still a significant downside risk for trust fund balances, which would lead to benefit cuts for future beneficiaries in certain states of the world, that downside is moderated when compared with current law projections in which the trust fund is almost certain to be exhausted.

Further exploration of the state-dependent policies provides more insights about how one might incorporate uncertainty into Social Security analysis and policymaking. If variation in fertility is ignored (that is, the overall fertility rate is always set to the expected value) then all of the downside risk for future benefits is eliminated relative to the base scenario of expected solvency after a one-time benefit cut. This is intuitive; the biggest long-run risk facing future beneficiaries in a partial pay-as-you-go system is that there might not be enough taxpayers to support them. This is the type of situation that might lead to a more fundamental change in the structure or scope of Social Security, meaning either higher payroll taxes to maintain the same benefits, or a reduction in benefits with tax rates held constant.

The other insights from the stochastic analysis concern which policies would be consistent with avoiding either explosion or implosion of the Social Security trust fund. Even

---

where near-term retirees are forced to accept benefit cuts deeper than those faced by future retirees, which would also probably violate conditions for political consensus regarding reform.

with longevity indexing and productivity-adjusted COLAs, trust fund balances are very large in good budgetary states of the world, meaning that benefits have been cut too much at some point in time. That suggests empirically answering the following question: Assuming that the two state-dependent policies analyzed here are put into place, what is the most that benefits can be initially cut to avoid significant chance of an exploding trust fund? The answer is about 5 percent, which is less than half of the 11 percent cut needed to achieve expected solvency. The fact that trust funds still implode in bad states of the world, however, leads naturally to a second question: Given that the two state-dependent policies are put into place, what is the least that benefits can initially be cut to avoid significant chance of an imploding trust fund? That answer is about 20 percent.

In discussion of these state-dependent benefit policies, the range between the 5 percent initial benefit cut that avoids trust fund explosion and the 20 percent initial cut that avoids trust fund implosion is useful to keep in mind when contemplating cohorts that might support various proposed Social Security reforms. It is highly unlikely that near-term retirees would agree to the 20 percent cut that virtually guarantees system solvency for future generations, but the criterion of balanced distributional outcomes across birth cohorts suggests that such a large cut would probably be inappropriate in any case. The modest 5 percent initial cut for near-term retirees seems both more acceptable and more distributionally neutral than current law.

2. The Relationship between Distributional Outcomes and Sustainable Solvency

The analysis of Social Security policy in this paper involves two related criteria: the implications for annual system financial flows and the ultimate distribution of taxes and benefits across birth cohorts. Comparing the distributional implications of any two policies on a level playing field requires holding the conditions for aggregate system financing fixed and requires working with time-invariant measures for distributional outcomes. In the analysis below, the condition for aggregate system financing is based on Lee and Yamagata's (2003) Flat Fund Ratio sustainable solvency criterion.[6] The distributional outcomes presented here are based on the ratio of lifetime benefits received to lifetime taxes paid.

Sustainable Solvency in Social Security

Sustainable solvency means that the Social Security trust fund ratio (trust fund balance divided by benefits paid) does not rise or fall in the long run. Given steady growth in real wages, a system in which benefits rise over time in step with real wages, and a steady-state ratio of beneficiaries to taxpayers, there is a wide range of policies that would all lead to sustainable solvency. Fundamentally, the larger is the trust fund in any given year, the less payroll tax revenue is required to pay any given level of current-year benefits and maintain a constant trust

[6]The goal of sustainable solvency could in principle be achieved using either benefit cuts or tax increases; Lee and Yamagata focused on payroll tax increases, but this paper restricts analysis to the benefit side. Benefit changes are the only policy lever considered here, in order to keep the scope more manageable for addressing the second policy issue, which is the distribution of the costs of achieving any given aggregate funding target. The distributional results presented here show how available resources under current law (taxes paid and interest on the trust fund) will be allocated across generations that are already alive, with the sustainable solvency condition that the trust fund not be deteriorating at the end of the simulation. The analysis here does not consider the possibility of imposing taxes on future generations whose lifetime tax and benefit outcomes are not being fully measured, mainly because that would be the distributional equivalent of a free lunch. This also differentiates the approach in this paper from the traditional risk-sharing analysis of Social Security as originally described by Gordon and Varian (1988) and more recently analyzed by Bohn (2001). That literature generally abstracts from trust fund financing, and thus taxes paid at any point in time equal benefits received.

fund ratio. Indeed, if the trust fund balance is sufficiently large, the interest alone could be enough to pay benefits and grow the fund.

Social Security faces two distinct types of aggregate financing problems. The first is associated with the aging of the baby-boom generation, which will cause a fundamental shift in the relationship between revenues and outlays lasting for the next several decades. The second source of financing problems is more long-run and systematic. If, as expected, mortality rates continue to fall and the fertility rate stays at recent levels, the ratio of beneficiaries to taxpayers will continue to rise even after the baby-boom generation has died off.

The steady-state situation described above is fairly close to what the U.S. is expected to face after the baby-boom generation has died off, and the only trend divergence between long-run system outlays and revenues is associated with increased longevity. That is, the underlying projections for fertility, real wage growth, and other factors are such that the system should be on a steady-state growth path if benefits are simply adjusted for longevity. Abstracting from longevity increases and holding payroll tax rates fixed, the fraction of scheduled benefits that are payable in the long run will be fixed and will depend only on how much is left in the Social Security trust fund after the baby boomers have stopped collecting benefits.

The balance remaining in the Social Security trust fund after the baby-boom generation dies off is an important connection between system financial and distributional outcomes. The analysis in this paper captures the effect of trust fund balances over time, because any shortfall in the trust fund automatically results in benefit cuts that show up in the distributional analysis across birth cohorts.[7] If the trust fund is allowed to fall to zero, as expected in the middle of te 21st century, benefits payable are assumed to be limited to payroll taxes collected thereafter.

Measuring Distributional Outcomes

The distributional effects of Social Security tax and benefit policy can be characterized and measured in several ways. The various distributional outcomes are sometimes referred to as "money's worth" measures (Geanakoplos, Mitchell, and Zeldes, 1999). The distributional outcomes shown in this paper are ratios of lifetime benefits to lifetime taxes, where both numerator and denominator are measured on a present-value basis. The money's worth outcomes are calculated for annual birth cohorts from 1950 through 2000.[8]

There are two basic strategies for measuring money's worth. The first is to compute the internal rate of return on Social Security contributions, which is the value for the discount rate that equates the present values of taxes and benefits. That is the value of r for which:

$$\sum_{t=0}^{100} taxes_t (1+r)^{100-t} = \sum_{t=0}^{100} benefits_t (1+r)^{100-t}$$

where $taxes_t$ and $benefits_t$ are taxes paid and benefits received at time t, respectively, and 100 is the maximum life span for the calculations. One significant problem with using the internal rate

[7] Although (for convenience) the actual simulations assume all benefits are cut across the board when the trust fund runs out of money, one could interpret the results as a sequence of cuts in new benefit awards made several years in advance.

[8] The values presented are actually centered five-year averages around single-year birth cohorts. The micro-simulation model used to compute the distributional outcomes (see the Appendix for details) is also capable of computing money's worth measures within cohort (for example, across lifetime income groups). The policy changes considered in this paper are restricted to proportional within-cohort benefit changes, in order to isolate the relationship between aggregate system financing and distribution over time, so all of the within-cohort distributional outcomes will change proportionally. For a more general analysis of within-cohort distributional outcomes, see, for example, Gustman and Steinmeier (2001) or Harris and Sabelhaus (2005).

of return is that readers sometimes (mistakenly) compare the outcomes to market rates of return, and make inappropriate normative judgments about the rate of return to Social Security (Geanakoplos, et al.).

The second approach to measuring money's worth is closely related but involves fixing the value for the real discount rate and taking the ratio of the present value of benefits to the present value of taxes. That is:

$$pvratio = \frac{\sum_{t=0}^{100} benefits_t (1+r)^{100-t}}{\sum_{t=0}^{100} taxes_t (1+r)^{100-t}}$$

Given that this equation has the same basic structure as the equation for calculating the internal rate of return, one would suspect the same basic critique about normative judgments is applicable. Taking the value for the real discount rate as given,

readers may be tempted to view a ratio above 1 as good and a ratio below 1as bad. In fact, present-value benefit to tax ratios should be compared only to other ratios that use the same real discount rate (Harris and Sabelhaus, 2005).

The interpretation of and qualifications about present-value ratios are best clarified by going to the actual data. Figure 1 shows present-value benefit-to-tax ratios for single-year birth cohorts from 1950 through 2000. The present value ratio for the 1950 cohort is estimated at 72 percent using a real discount rate of 3.3 percent. That is, the present value of all lifetime benefits received (and projected to be received) by the 1950 cohort is about 72 percent of the present value of all lifetime taxes paid (and projected to be paid).[9]

There are two sets of projections for present-value ratios across birth cohorts in Figure 1, computed using two extreme assumptions about Social Security system financing. The higher of the two, Scheduled Benefits, shows what would be expected of the money's worth measure across cohorts if all promised benefits were somehow (unrealistically) paid with just scheduled tax receipts. The general upward trend (the present-value ratio for the 2000 cohort is about 10 percent higher than the ratio for the 1950 cohort) reflects the fact that future beneficiaries are expected to live longer and thus receive higher lifetime benefits.

In the alternative Current Law scenario, across-the-board benefit cuts are implemented when the Social Security trust fund is exhausted. Figure 2 shows that in the projections underlying these money's worth calculations, the Social Security trust fund is exhausted around the middle of the 21st century, so every cohort for which there are any surviving beneficiaries after that time will be affected. Of course, the longer a cohort receives benefits post-trust fund exhaustion, the bigger the lifetime benefit cuts. The projected present-value ratio for the 2000 cohort, expected to receive virtually all of their benefits after trust fund exhaustion, is more than 20 percent lower than the present-value ratio for the 1950 cohort.

---

[9]The calculations include all worker and auxiliary benefits paid under the Old-Age and Survivors Insurance and Disability Insurance programs, excluding only children's and young spouse beneficiaries. To underscore the sensitivity of the estimates to assumptions about real rate of return, Harris and Sabelhaus (2005) observe that setting the discount rate equal to 2.2 percent generates estimated ratios of around 100 percent. That sensitivity makes sense, given the differences in timing of taxes and benefits over the life cycle. The same analysis also shows, however, that the relative estimates of distributional outcomes across groups are not much affected by the discount rate. The value of 3.3 for the discount rate is used here (and in various Congressional Budget Office publications) in order to maintain consistency with the assumed rate of return on government bonds.

3. Uncertainty about System Finances and Distributional Outcomes

Analysis of Social Security policy along the two dimensions described above, aggregate system finances and distributional outcomes, requires an integrated modeling framework that generates both macroeconomic and microeconomic results.[10] Given the modeling framework, projections can be deterministic or stochastic. Stochastic simulation explicitly accounts for the uncertainty underlying the projections by combining many simulations, each based on random draws for every input assumption in each year of the projection. Any one simulation of this type has very limited usefulness, it can be thought of only as one possible realization. However, the collection of multiple simulation results represents a probability distribution for the system financial and distributional outcomes of interest. That makes it possible to investigate state-dependent policy rules such as those being considered here.

The model used for this paper has 11 stochastic inputs: fertility rate, rate of mortality improvement, rate of immigration, disability incidence, disability termination, total factor productivity growth, earnings share of compensation, unemployment rate, inflation rate, gap between the real interest rate and the marginal product of capital, and gap between the core GDP deflator and inflation.[11] For each input, the underlying equations reflect historical variability and covariances, and these equations, together with random draws for the equation error terms, are used to generate annual values.

The benefit of moving to stochastic simulation is shown in Figure 3, which presents the uncertainty bands around the expected ratios of lifetime benefits to lifetime taxes for the 1950

---

[10]See the Appendix and references therein for a complete description of the model used here.

[11]For a detailed discussion of the stochastic model and underlying time series processes, see Congressional Budget Office (November 2005).

through 2000 birth cohorts under the same two scenarios (Scheduled Benefits and Current Law) presented in Figure 1. The upper and lower bands represent the 10^{th} and 90^{th} percentiles of the estimated probability distribution for the present value ratios. That is, there is a 10 percent chance that the ratio of lifetime benefit to lifetime taxes will be below the lower band and a 10 percent chance that the ratio will be above the upper band.

The basic shapes of the ranges for present value ratios are consistent with the deterministic projections shown in Figure 1, but Figure 3 shows how two different sources of uncertainty affect the ultimate outcomes. The range for Scheduled Benefits (the cross-hatched band) reflects the fact that uncertainty in demographic and economic assumptions makes the ratio of lifetime benefits to taxes uncertain, and that ratio's uncertainty increases the further ahead one looks. The range for near-term retirees is narrower than for distant retirees because lifetime taxes paid are already largely known. Benefits are also very predictable for near-term retirees because they are based on mostly known earnings histories. There is still uncertainty about mortality rates, inflation, and other variables, but the range for lifetime benefits to taxes is much narrower than for retirees in the distant future, who face an entire lifetime of uncertainty.

The range for Current Law (the dark shaded band) includes all of the uncertainty associated with Scheduled Benefits, but adds uncertainty about trust fund exhaustion. Recall that when the trust fund is exhausted, scheduled benefits can no longer be paid, and all beneficiaries experience a benefit cut such that benefits paid are equal to payroll tax revenues in that year. The fact that the 90^{th} percentile for Current Law benefits is below the 90^{th} percentile for Scheduled Benefits indicates that the probability of trust fund exhaustion is above 90 percent at some point during that cohort's life span. The 10^{th} percentile for Current Law benefits lies well below the 10th percentile for Scheduled Benefits. The trust fund is very likely to become exhausted and thus force benefit cuts, and there is a potential for those cuts to be quite large.

The stochastic relationship between system financial and distributional outcomes is driven home by considering the projected range for trust fund ratios, shown in Figure 4. The projections clearly show not onlythat the Social Security system is on an expected path to insolvency, but also that the probability that it maintains solvency over the next 100 years is negligible. At the 10th percentile, the trust fund is exhausted by 2035, suggesting that there is a 10 percent chance of depletion occurring before 2035. Even at the 90th percentile, the trust fund is exhausted by 2088 under Current Law, indicating that 90 percent of the time scheduled benefits are no longer payable by 2088. It is this near-certainty about trust fund exhaustion that drives such a large wedge between the ranges for distributional outcomes shown in Figure 3.

4. Achieving Sustainable Solvency with a One-Time Benefit Cut

There are many approaches policymakers could take to resolve the expected long-run financing shortfall in Social Security. This section considers a one-time permanent benefit cut that is expected to lead to the Flat Fund Ratio sustainable solvency described earlier. The strategy is to impose a fixed percentage reduction (relative to scheduled benefits) for everyone 55 and younger as of 2005. The benefit cut used in the simulation, 19 percent, was chosen to meet the sustainable solvency criterion as closely as possible with a single cut in a single year.

Figure 5 shows that the expected sustainable solvency outcome is achieved by using the 19 percent cut. The dramatic decrease in benefits leads to faster trust fund accumulation starting in 2012, when the first reductions are implemented. The eventual leveling out of the trust fund ratio, which occurs after the baby boom generation is expected to have died off, is at a level nearly twice as high as the peak (around 2020) under current law. Thus, the simulation can be interpreted as follows: If all retirees claiming benefits after 2012 accept a large cut relative to current law, that enables the trust fund to stay solvent and reach an equilibrium level where payroll taxes plus interest earned by the fund are sufficient to pay benefits and maintain a stable trust fund ratio.

The effect of this policy on money's worth ratios across cohorts (shown in Figure 6) is predictable, causing significant redistribution from near- and medium-term retirees to future retirees. There are two reinforcing reasons. First, money's worth ratios in the 19 percent Benefit Cut scenario show the same pattern (though at a lower level) as seen in the Scheduled Benefit scenario in Figure 1. That is, relative to near-term retirees, money's worth ratios rise in the future because of increased longevity. Second, because the trust fund is not being exhausted benefits do not need to be reduced as under current law. Therefore, relative to current law, generations born after the early 1980s would find the 19 percent benefit cut to be preferable, but those born before would be worse off.

The effect of a large one-time benefit cut can also be evaluated using the stochastic analysis. Figure 7 shows the uncertainty bands for money's worth ratios under current law (reproduced from Figure 3) and the 19 percent cut. The inter-generational redistribution apparent in Figure 6 shows up in Figure 7 through a general downward shift in the uncertainty bands for near and medium-term retirees. One insight about how stochastic analysis contributes to understanding of the political aspect of the problem comes from comparing the points where various lines cross in Figures 6 and 7. In a deterministic world (Figure 6), groups born before the early 1980s are better off sticking with current law (money's worth ratios are higher). In a stochastic world, if people are more concerned about worst-case scenarios (say at the 10^{th} percentile), then the break-even point shifts back to about the 1970 birth cohort.

In addition to the expected shift in resources across generations, the range of money's worth ratios for the 19 percent cut is much narrower in the long run, which is consistent with the stochastic projections for Social Security trust fund ratios shown in Figure 8. The upside potential for money's worth is mitigated because everyone faces the 19 percent benefit cut in all states of the world. This is worse than the best possible outcome under Current Law, because achieving expected solvency means that sometimes benefits will have been cut too much, as reflected in an exploding trust fund ratio at the 90^{th} percentile (Figure 8). The downside potential for money's worth is mitigated because the risk of trust fund exhaustion is greatly reduced.

5. Stochastic Outcomes Using State-Dependent Policy Rules

The 19 percent one-time Benefit Cut scenario analyzed in the previous section achieves expected Social Security system solvency, but the distributional consequences are rather extreme, especially for near-term retirees. It is not obvious what normative reasoning could justify a policy that lowers expected money's worth ratios more for near-term retirees than it does for future retirees. One could argue the opposite that, because future retirees are expected to enjoy higher real living standards, giving them a lower rate of return on their taxes paid might be appropriate. In any case, the policy leaves a lot of money on the table in a stochastic sense. The exploding 90th percentile for the Social Security trust fund ratio in Figure 8 shows that there is a good chance that benefits will have been cut more than needed to maintain a stable trust fund ratio in the long run.

The analysis in this paper suggests that introducing state-dependent policies could help balance the goals of achieving certain distributional objectives and maintaining system solvency.[12] In particular, if future benefit cuts are tied to realized demographic and economic outcomes, one could start with smaller initial benefit cuts for near-term retirees, tie benefits to economic and demographic outcomes in a way that directly affects money's worth calculations, and recapture the upside potential for future retirees in good states of the world by deferring benefit cuts when they are not needed. Although the underlying uncertainty about demographic and economic outcomes cannot be changed, that uncertainty can be systematically incorporated into policymaking.

---

[12]It is worth noting that other distributional objectives would not necessarily be consistent with the strategy of putting the system on "auto-pilot" using state-dependent policy rules. One example is the objective of making sure that benefit levels for the lowest earning retirees are generally sufficient to alleviate elderly poverty.

Longevity Indexing and Productivity-Adjusted COLAs

There are two state-dependent policies analyzed here. The first ties future benefit cuts to changes in longevity outcomes over time. That would shift a larger portion of the long-term funding burden onto future retirees, easing the burden on near-term retirees, but in such a way as to automatically flatten the money's worth ratios across cohorts. In a deterministic simulation, the sustainable solvency outcome is achieved using an 11 percent initial benefit cut in 2012 plus longevity indexing.[13]

The second state-dependent rule analyzed here ties post-retirement benefit adjustments (also known as COLAs) to productivity and Consumer Price Index (CPI) inflation, rather than just to CPI inflation as under current law. If productivity growth is strong, the Social Security system becomes more solvent because payroll taxes rise immediately but benefits increase only after a lag. Indeed, a large part of the upside potential for the trust fund ratio shown in Figure 8 is because of the upside potential for productivity growth. With productivity-adjusted COLAs, benefits paid would be immediately higher when productivity growth is strong, which shifts some of that upside potential for the trust fund to higher money's worth for retirees. However, in order to avoid *expected* benefit increases, the proposal simulated here involves starting with CPI inflation, then adding the *deviation* in productivity from its average (1.2 percent per year) to determine the overall COLA. In that sense, below-average productivity outcomes would lead to a decline in benefits relative to the inflation-only COLA.[14]

---

[13]The version of longevity indexing used here is the same as that proposed in The Bipartisan Retirement Security Act of 2004 (H.R. 3821), better known as the Kolbe-Stenholm reform plan.

[14]The analysis here abstracts from normative judgment about how benefits should change after claiming begins. Feldstein (1990) considers arguments for whether benefits should be tied to inflation or grow with productivity. Also, an interesting aspect of the Diamond and Orszag (2004) Social Security reform plan is the

Figure 9 compares money's worth ratios under the one-time 19 percent benefit cut (reproduced from Figure 5) to the results from the 11 percent cut along with the two state-dependent policy rules. Again, both of these policy packages exhibit expected sustainable solvency, so they are comparable. The comparison clearly shows a stochastic improvement in the upside potential for money's worth outcomes for both near-term and future retirees, because the initial benefit cut is lower and future benefit changes are tied to longevity and productivity outcomes. Indeed, the 90^{th} percentile for money's worth ratios is nearly flat across birth cohorts. As before, the pattern of money's worth outcomes is tied directly to trust fund outcomes, as shown in Figure 10. The upside potential for money's worth ratios is higher under the state-dependent policy rules because the top of the range for trust fund ratios is much lower.

Figure 9 also shows that there is some deterioration relative to the 19 percent cut at the low end (10^{th} percentile) for future retirees. That occurs because the probability of trust fund exhaustion is somewhat higher under the state-dependent policies than it is under the 19 percent benefit cut, as shown in Figure 10. Still, the difference in downside potential is fairly modest relative to the overall range for money's worth outcomes. Also, the 10^{th} percentile of money's worth outcomes for future retirees is still well above the downside outcome under current law (Figure 7).

"super-COLA" for disabled worker beneficiaries, which would grow benefits at a rate above CPI inflation in order to avoid the lifetime (relative) deterioration that occurs under current law.

The Role of Uncertain Fertility

Although Figure 10 exhibits considerably less upside potential for the trust fund ratio under the state-dependent policy rules plus the 11 percent cut than under the 19 percent cut, other uncertain inputs still affect trust fund outcomes. The 90th percentile for the state-dependent policy rules plus the 11 percent cut shows the same upward trend as the 90th percentile for the 19 percent cut (an exploding trust fund) although at a somewhat slower pace. Also, the downside risk for system solvency is now somewhat worse, because starting with only an 11 percent initial benefit cut increases the probability of exhaustion relative to the 19 percent cut. Adjusting future benefits for longevity and productivity are not powerful enough to reverse that effect.

Fertility rates rank among the most important of the other determinants of uncertainty about trust fund balances.[15] It is not immediately obvious how one could tie benefit changes to fertility rates, but a simple experiment shows how important fertility variation is for future system solvency and money's worth outcomes. Figures 11 and 12 show two versions of money's worth ratios and trust fund outcomes under the 11 percent benefit cut plus state-dependent policy option. The first set is reproduced from Figures 9 and 10, with all sources of uncertainty considered. The second set of outcomes are for the same policy simulations, but with the stochastic variation in fertility turned off.

The most noticeable effect of shutting down the variation in fertility rates is that the downside risk for money's worth outcomes for future retirees is greatly mitigated. Indeed, the 10th percentile for future retirees is now roughly equal to that for the 19 percent benefit cuts. A major risk facing future retirees in a partially pay-as-you-go system is that there will not be

[15] This observation is based on the decomposition of overall uncertainty about trust fund outcomes into its component sources, as shown in Congressional Budget Office (November 2005).

enough future taxpayers to fund their benefits. The upside potential for money's worth ratios is not affected by shutting down the fertility variation, because having more than enough taxpayers to fund benefits does not lead to an increase in benefits. It leads instead to an increase in trust fund balances.

This powerful relationship between fertility, trust fund ratios, and money's worth outcomes suggests that there is a higher level of state-dependent policymaking that could come into play in certain scenarios. In particular, Sheshinski and Weiss (1981) show that, in an overlapping-generations model with uncertain life span, the optimal level of Social Security rises if fertility rates fall. That suggests future retirees might decide, if they choose to have fewer children than currently projected, to pay more in taxes while working in order to fund their own benefits. The alternative to paying the higher taxes is to receive lower benefits.

Avoiding Exploding and Imploding Trust Funds

Adopting state-dependent policy rules for determining Social Security benefit changes over time is an important step toward balancing the objectives of stable system finances and fairness in distributional outcomes. The two state-dependent rules considered here move only part of the way towards a system that is robust with respect to demographic and economic outcomes, and there are no other obvious trigger mechanisms that could be used to offset variation in the other inputs. It is useful to explore the issue from another point of view, however: to consider what range of initial benefit changes, when combined with the state-dependent policies, is consistent with trust fund ratios that neither explode nor implode.

Even with longevity indexing and productivity-adjusted COLAs, trust fund balances are very large in good budgetary states of the world, meaning that benefits have been cut too much at some point in time. The trust fund explodes in good states of the world because the benefit cuts put in place into 2012 are (ex post) too deep if outcomes are favorable to system finances. So, the first question to ask is this: Assuming that (1) any benefit cuts put in place will never be reversed and (2) benefits will be adjusted for longevity and the COLA will be adjusted for productivity, what is the most that benefits could be reduced in 2012 without introducing the possibility of exploding trust fund ratios?

There is no direct way to solve for the maximum initial cut that rules out an exploding trust fund. The only feasible approach is to choose a value, use the stochastic simulation capability, and observe what happens to the 90th percentile of the trust fund ratios. The answer generated by this search algorithm is 5 percent. If benefits are initially cut by 5 percent in 2012 and then are indexed to longevity and adjusted for productivity, the 90th percentile of the trust fund ratio is basically stable after the middle of the century. Thus, it can be argued that an initial cut of 5 percent, when combined with the state-dependent policy rules, would rarely lead to a situation in which, in review, benefits had been cut too much for near-term retirees.

The 5 percent cut that avoids trust fund explosion is less than half the 11 percent cut needed to achieve expected solvency, and thus the probability of an imploding trust fund is higher. That leads naturally to a second question: Given that the two state-dependent policies are put into place, what is the least that benefits can be cut initially while still avoiding a high probability of an imploding trust fund? That is, is there a policy such that the 10th percentile of the trust fund ratio range flattens out and stays parallel with the 90th percentile into the indefinite future? This answer, also solved for by searching over possible values for the initial cut, is an initial benefit cut (in 2012) of about 20 percent. The wide range between the initial changes that prevent trust fund explosion or implosion arises in part because of inherent uncertainty, but also because of the underlying demographics expected to drive trust fund balances. Basically, only a huge cut for near-term retirees could guarantee a solvent system for the indefinite future.

6. Conclusions

The Social Security system is expected to experience financing problems sometime in the next few decades, but the exact timing and magnitude of those problems are uncertain. In such a situation, systematically tying policy rules to realized economic and demographic trends could help balance the objectives of maintaining system solvency and targeting distributional outcomes. Although uncertainty about economic and demographic trends cannot be eliminated, it can be addressed, as reflected in a general narrowing of the size and change in the shape of uncertainty bands around trust fund ratios over time and money's worth outcomes across cohorts.

One of the implications of this analysis has to do with the impetus for changing Social Security. Near-term retirees may have little interest in participating in a reform discussion, because the system is not expected to experience funding shortfalls for several decades. Future retirees, on the other hand, may begin to realize that failure to implement any benefit changes for near-term retirees means they will be bearing all the costs of shortfalls under current law because their benefits will be reduced after the Social Security trust fund is exhausted. Their support for the existing system could wane.

If the system's goals are sustainable solvency and reasonably stable money's worth ratios, then a conditional approach to policy, a combination of longevity indexing and productivity-adjusted benefits, could alter the prospects for change. The analysis here shows how a fairly modest initial benefit cut followed by further changes tied to economic and demographic outcomes would lower the chance of future system insolvency while maintaining a relative narrow range for distributional outcomes.

Appendix: The CBOLT Micro-Simulation Model

The Congressional Budget Office Long-Term (CBOLT) policy analysis tool integrates microsimulation with a long-run macroeconomic and federal budget simulation model. The microsimulation component of CBOLT has two phases. The first phase involves generating realistic demographic and economic outcomes for a large representative sample of the population. The second phase involves applying Social Security program rules to determine individual taxes and benefits. The taxes and benefits solved for in the microsimulation are aggregated to solve for overall system financial outcomes, and are also used for comparative analysis of outcomes across various demographic and socioeconomic groups. One feature of the CBOLT macroeconomic framework also plays an important role in the analysis here: The model can be solved repeatedly, using random draws for key economic and demographic assumptions, which allows presentation of outcomes in terms of probability distributions, rather than just in expected values.

The starting point in any microsimulation model is the base data file. CBOLT uses the Social Security Administration (SSA) Continuous Work History Sample (CWHS), which is a 1-in-100 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 1-in-1000 sample), the data are 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 microcalibration process discussed below.

Given a micro database file, the next step in any microsimulation is to specify transition equations, the processes by which demographic and economic outcomes for individuals in the microsample 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 death) needed to calculate Social Security taxes and benefits and to integrate the microeconomic outcomes with the macroeconomic growth model and unified budget framework. The micro transition processes are described in detail in a series of technical papers available on the Congressional Budget Office web site. For a more detailed overview of the micro-modules, see O'Harra, Sabelhaus, and Simpson (2004).

In general, the CBOLT microtransitions equations are not designed to incorporate dynamic optimizing behavior. Rather, the focus is on identifying and replicating real-world microheterogeneity in the simulated sample. However, even this limited-behavior version of the microsimulation-based approach leads to important insights about Social Security that do not come through in other analyses. First, all else equal, the microsimulation 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 (Congressional Budget Office, 2004). 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.

The microtransition 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-database file. Although this type of imputation-based assignment is used in all major microsimulation projects, it is quite extensive in CBOLT and deserves special mention. The basic idea is to assign missing characteristics in history by using the same methods used for projecting forward, then to test and calibrate the processes by using external data sources available in history.

The initial CBOLT microsample is actually drawn to be representative of the U.S. population for the period 1984 through 2003. Social Security coverage rates were much lower for working-age cohorts before then; also, the availability of earnings data above the taxable maximum did not occur until the 1980s. The historical simulation begins in 1984, when each individual is assigned initial unobserved characteristics based on observed characteristics (a standard imputation). Next, each micro-process is applied and then calibrated so that 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.

Uncertainty is explicitly accounted for in CBOLT by using a stochastic macrodemographic model that generates values for certain demographic and economic inputs. The stochastic modules in CBOLT are related to other models of this type, including Holmer

(1999), Lee and Tuljapurkar (1998), and Lee and Edwards (2002). The key demographic input assumptions varied stochastically in the Monte Carlo simulations are fertility, mortality, net immigration, and rates of disability incidence and termination. The key economic inputs varied stochastically are total factor productivity growth, inflation, unemployment, the relationship between interest rates and the return to capital, the share of compensation that shows up as taxable earnings, and the gap between the core GDP deflator and growth in the consumer price index for urban wage earners and clerical workers (CPI-W). These factors together determine the individual and aggregate outcomes of interest in any simulation (for further details see Congressional Budget Office, November 2005). In this paper, all stochastic results are based on 500 simulations.

References

Bohn, Henning (2001). "Retirement Saving in an Aging Society: A Case for Innovative Government Debt Management." University of California, Santa Barbara, Economics Working Paper Series, Number 1056.

Congressional Budget Office (June 2004). *The Outlook for Social Security*. Washington, D.C.

Congressional Budget Office (November 2005). *Uncertainty in the Analysis of Long-Term Social Security Projections.* Washington, D.C.

Congressional Budget Office (March 2005). *Updated Long-Term Projections for Social Security*. Washington, D.C.

Cornia, Gary C., Ray D. Nelson, and Andrea Wilko (2004). "Fiscal Planning, Budgeting, and Rebudgeting Using Revenue Semaphores." *Public Administration Review*, Vol. 64, No. 2 (March/April), pp. 164-179.

Diamond, Peter A. and Peter R. Orszag (2004). *Saving Social Security*. Washington, D.C.: Brookings Institution Press.

Feldstein, Martin (1990). "Imperfect Annuity Markets, Unintended Bequests, and the Optimal Age Structure of Social Security Benefits." Journal of Public Economics, 41:1, 31-43 (February).

Geanakoplos, John, Olivia S. Mitchell, and Stephen P. Zeldes (1999). "Social Security Money's Worth." *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.

Gordon, Roger H. and Hal R. Varian (1988). "Intergenerational Risk Sharing." *Journal of Public Economics*, Vol. 37, pp. 185-202.

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

Harris, Amy Rehder and John Sabelhaus (May 2005) "How Does Differential Mortality Affect Social Security Progressivity and Finances?" Technical Paper 2005-5. Congressional Budget Office, Washington, D.C.

Harris, Amy Rehder and Michael Simpson (2005) "Winners and Losers Under Various Approaches to Slowing Social Security Benefit Growth," *National Tax Journal*, Vol. 58, No. 3 (September), pp. 523-544.

Holmer, Martin R. (1999). “Stochastic Simulation of Economic Growth Effects of Social Security Reform.” *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.

Lee, Ronald and Ryan Edwards (2002). “The Fiscal Impact of Population Aging in the U.S.: Assessing the Uncertainties.” Center for Economics and Demography of Aging, University of California, Berkeley.

Lee, Ronald and Shripad Tuljapurkar (1998). “Stochastic Forecasts for Social Security.” in *Frontiers in the Economic of Aging*, ed. David A. Wise, Chicago and London: University of Chicago Press.

Lee, Ronald and Hisashi Yamagata (2003). “Sustainable Social Security: What Would It Cost?” *National Tax Journal*, Vol. 57, No.1, Part 1 (March) pp. 27-43.

O’Harra, Josh, John Sabelhaus, and Michael Simpson (January 2004) “Overview of the Congressional Budget Office Long-Term (CBOLT) Policy Simulation Model,” Technical Paper 2004-1. Congressional Budget Office, Washington, D.C.

Popper, Steven W., Robert J. Lempert and Steven C. Bankes (2005). “Shaping the Future: Overview/Dealing with Uncertainty.” *Science*, March 28.

Sheshinski, Eytan, and Yoram Weiss (1981). “Uncertainty and Optimal Social Security Systems.” *Quarterly Journal of Economics*, 96:2, 189-206 (May).

Smetters, Kent (2004). “Solvency and Reform of Social Security: Is the Social Security Trust Fund a Store of Value?” *American Economic Review*, Vol. 94, No. 2 (May), pp. 176-181.

Social Security Administration Board of Trustees (April 2005). “The 2005 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds.” Available at www.ssa.gov/OACT/TR/TR05/tr05.pdf.

Figure 1. Social Security Benefit-to-Tax Ratios by Single-Year Birth Cohorts

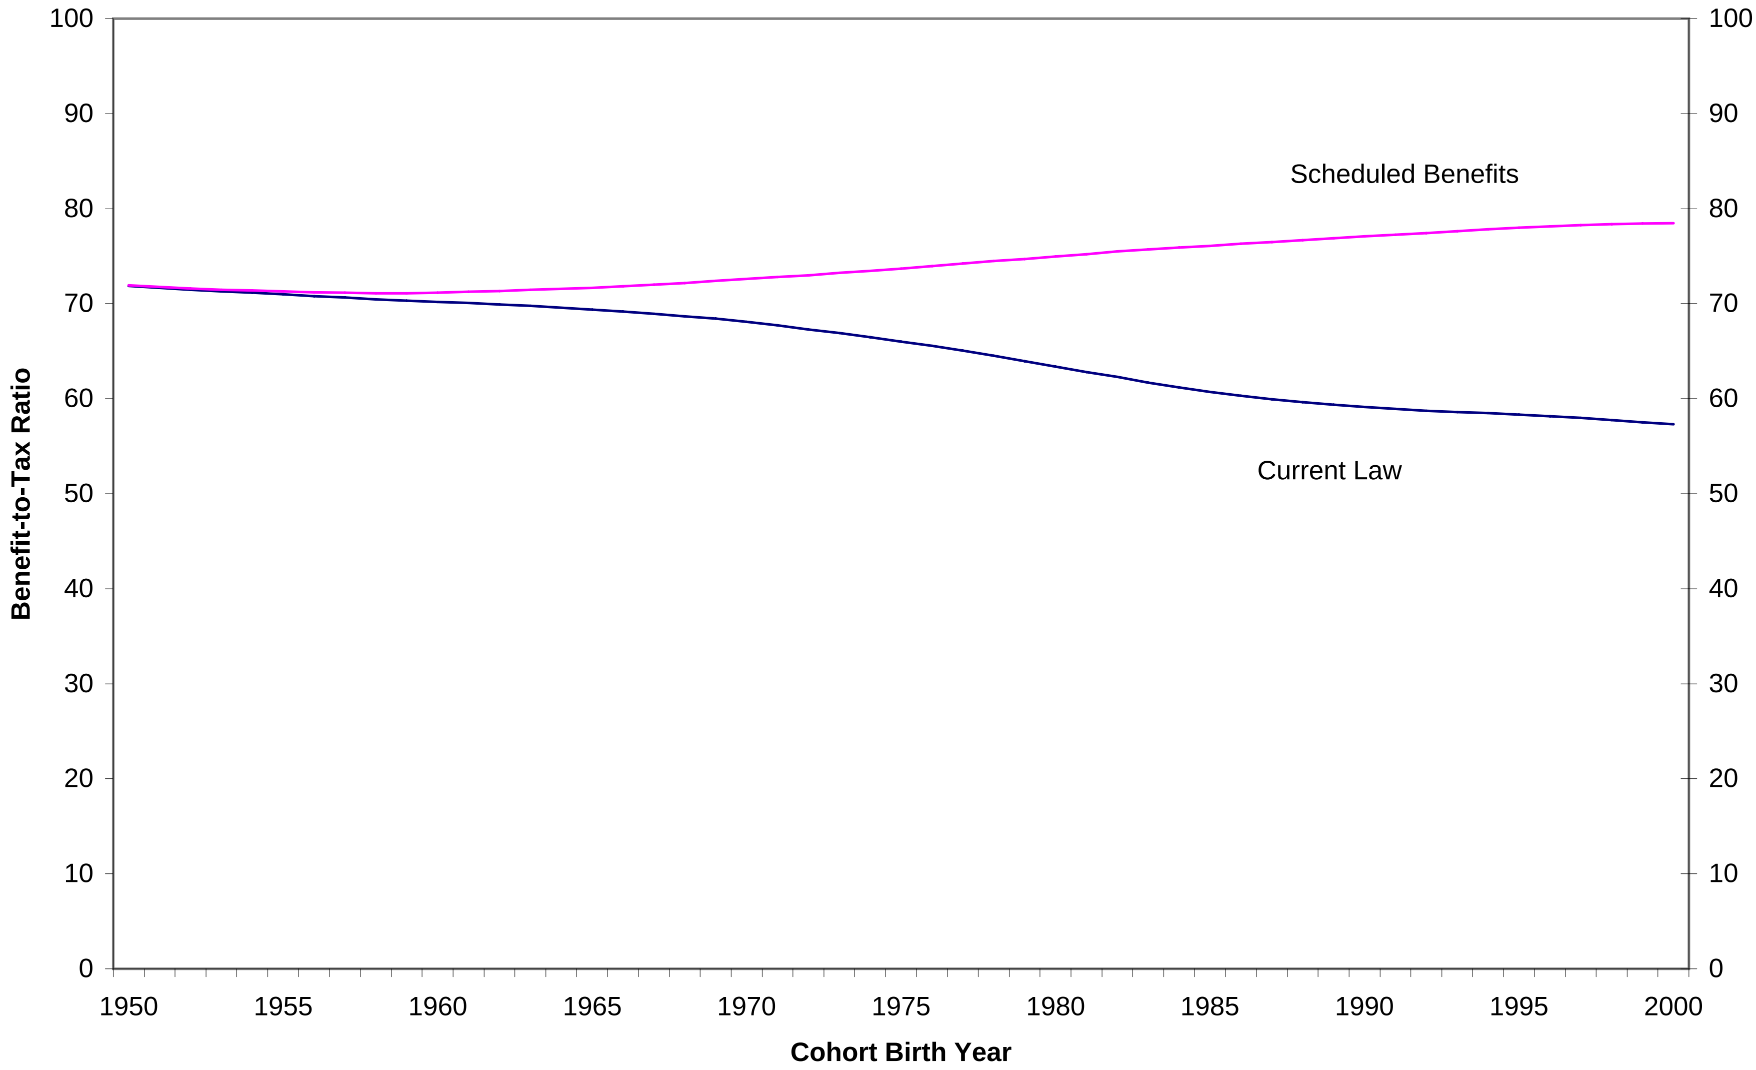

Figure 2. Expected Social Security Trust Fund Ratio

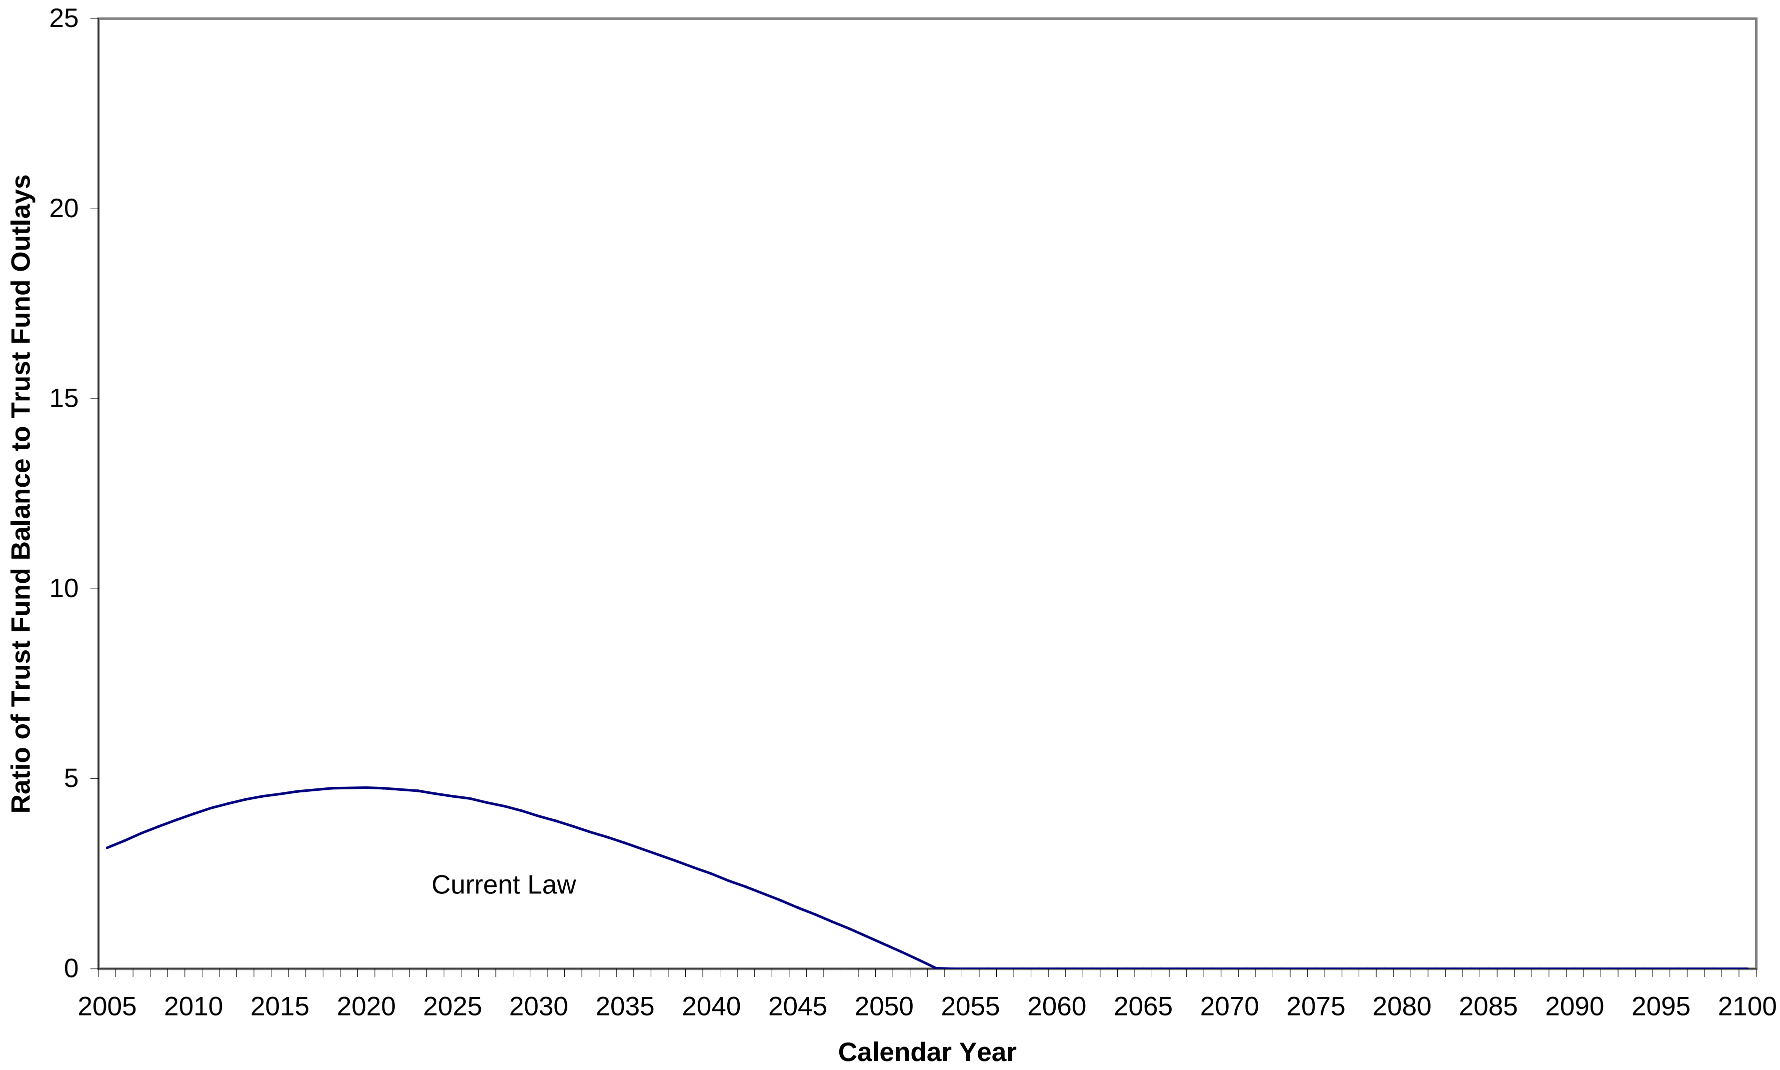

Figure 3. Potential Ranges for Ratios of Lifetime Benefits to Lifetime Taxes

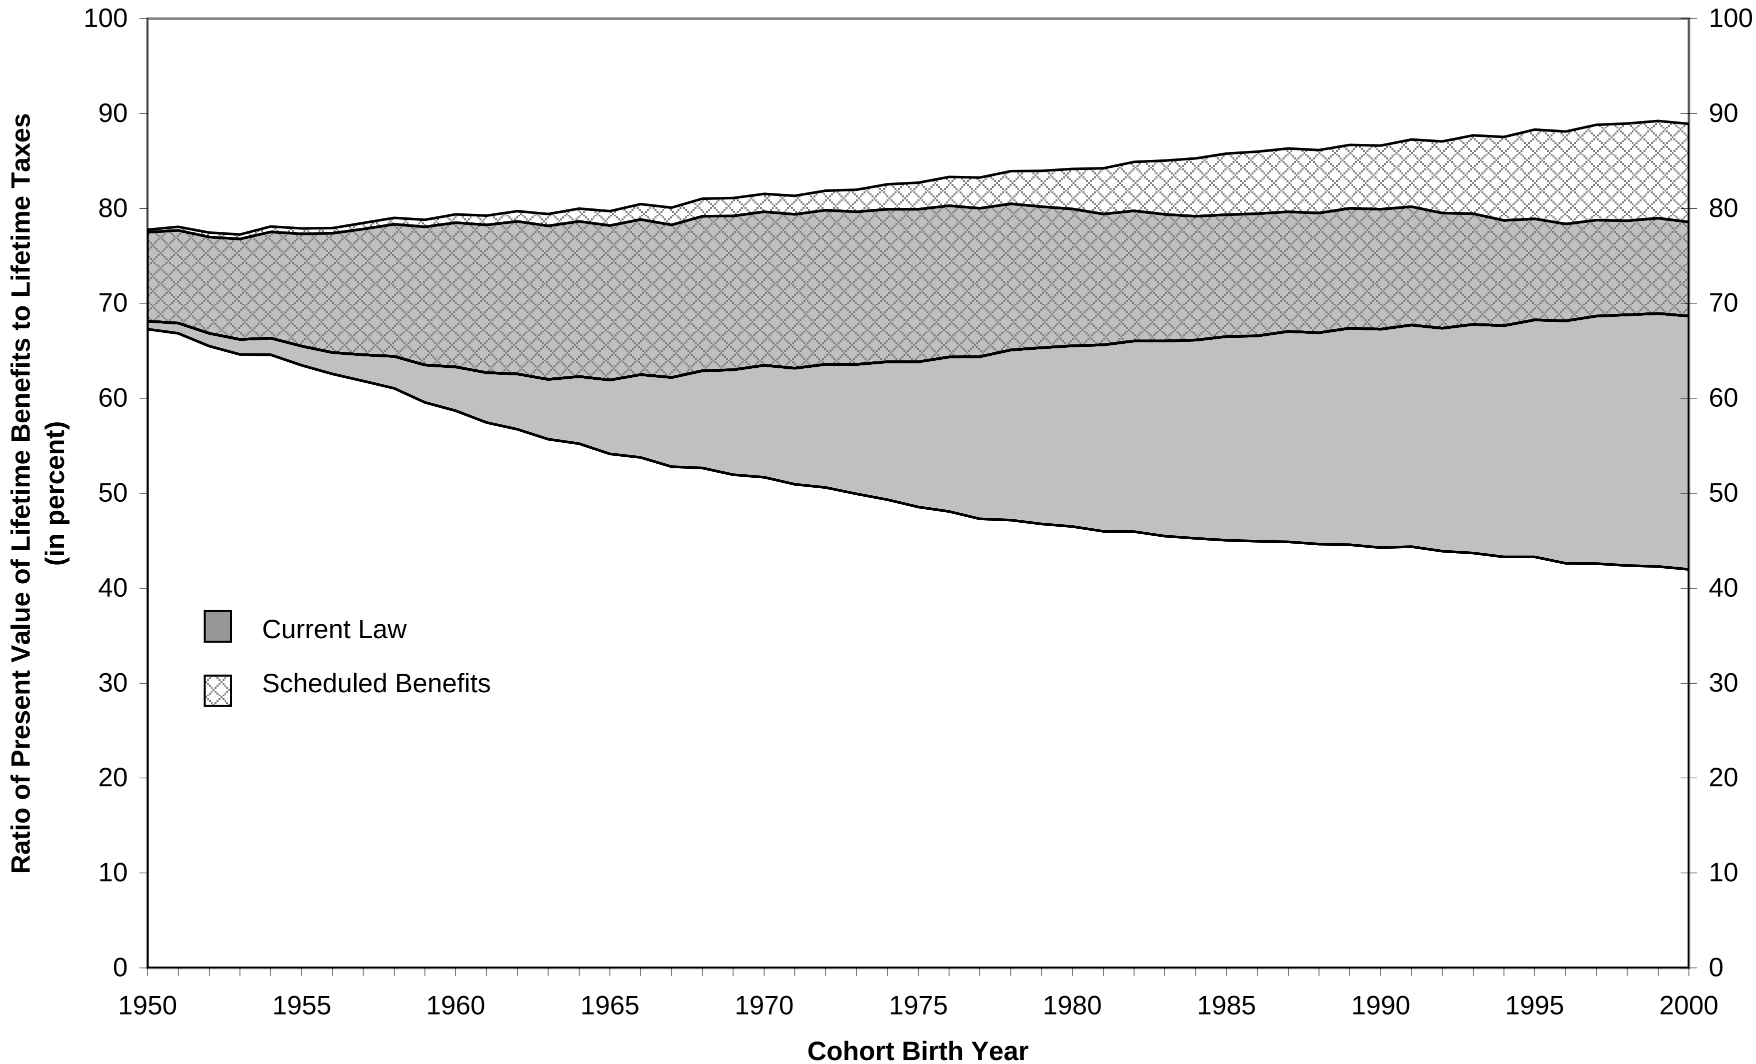

Figure 4. Potential Range for Social Security Trust Fund Ratios under Current Law

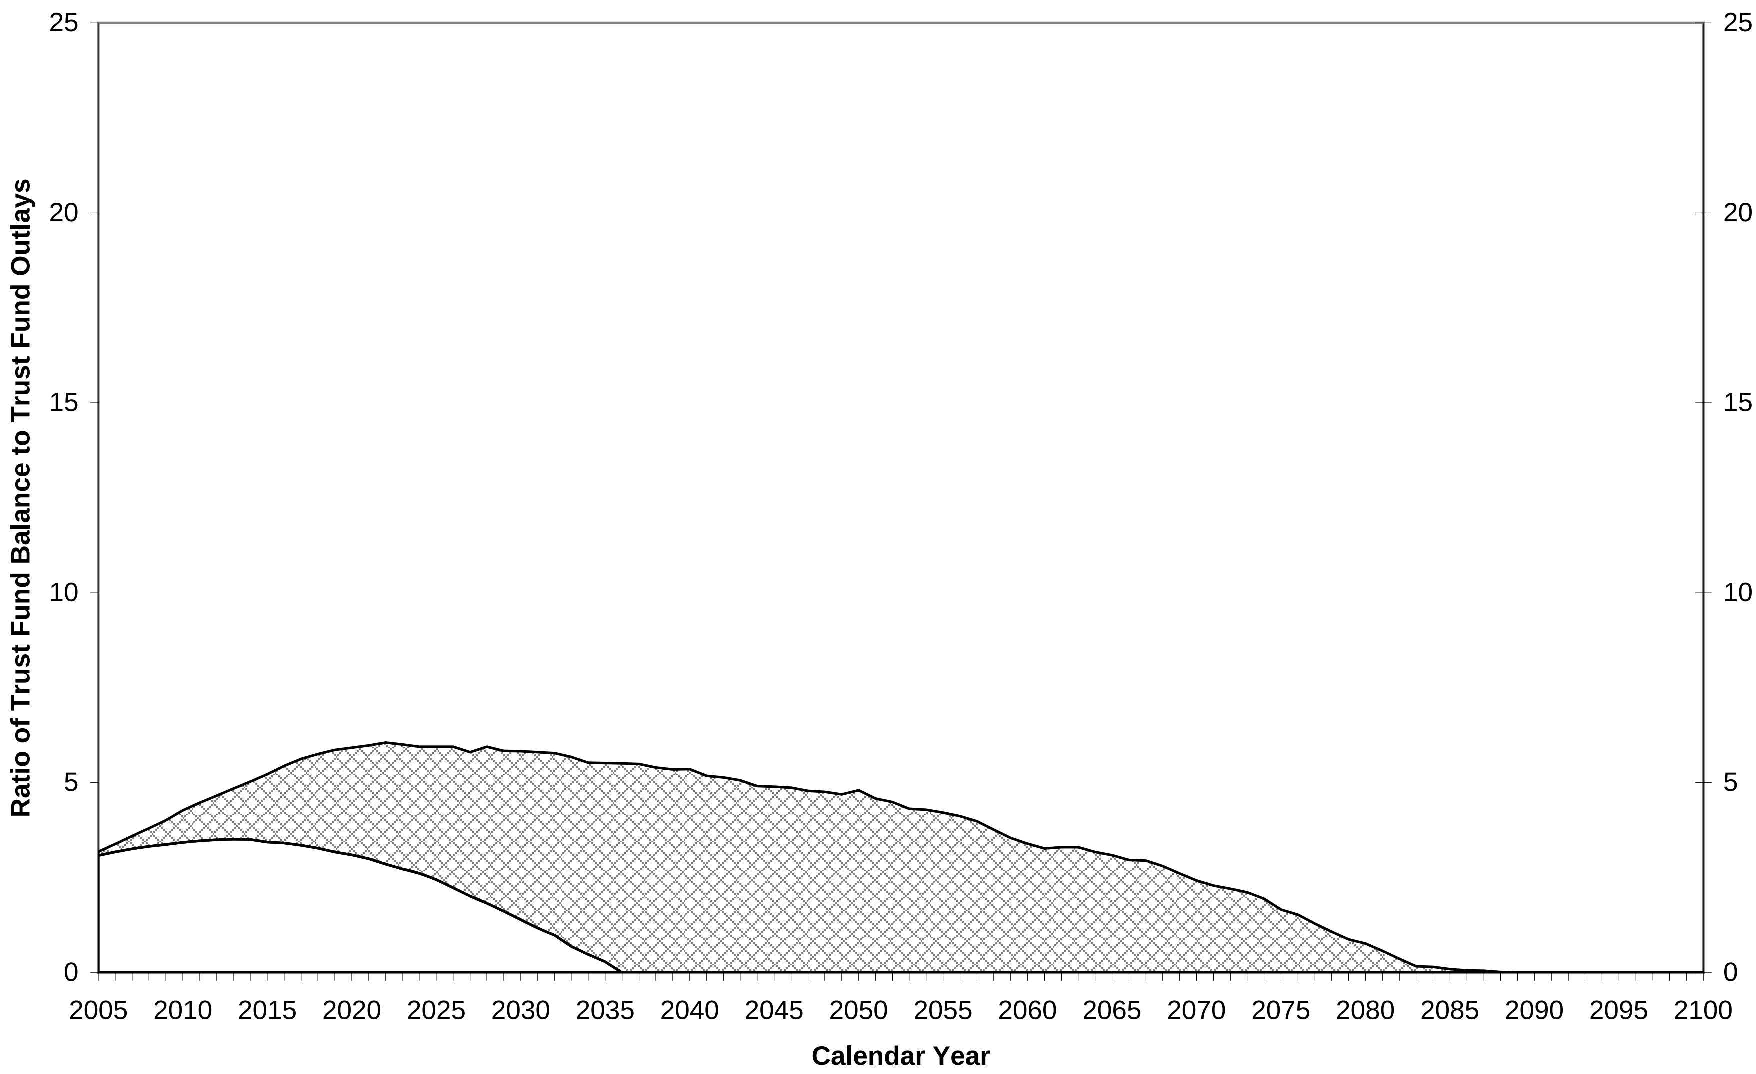

Figure 5. Social Security Trust Fund Ratios
(Current Law and One-Time Sustainably Solvent Benefit Cut)

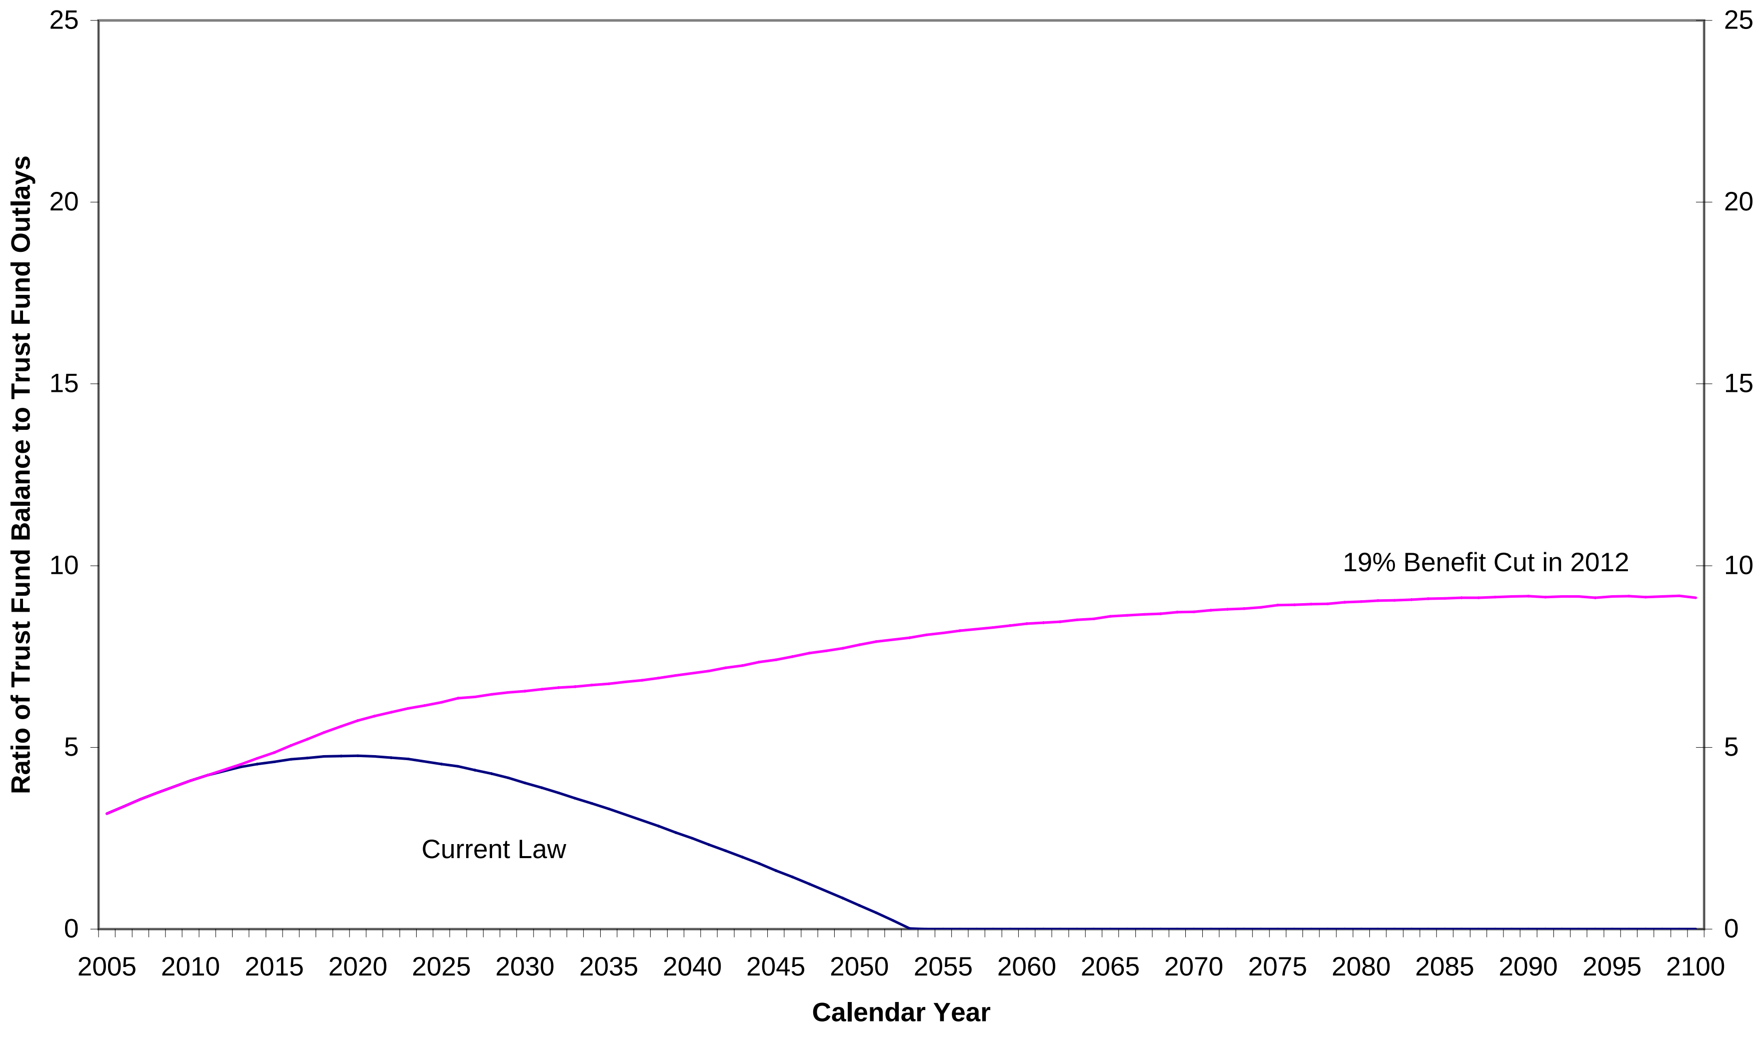

Figure 6. Ratios of Lifetime Benefits to Lifetime Taxes (Current Law and One-Time Sustainably Solvent Benefit Cut)

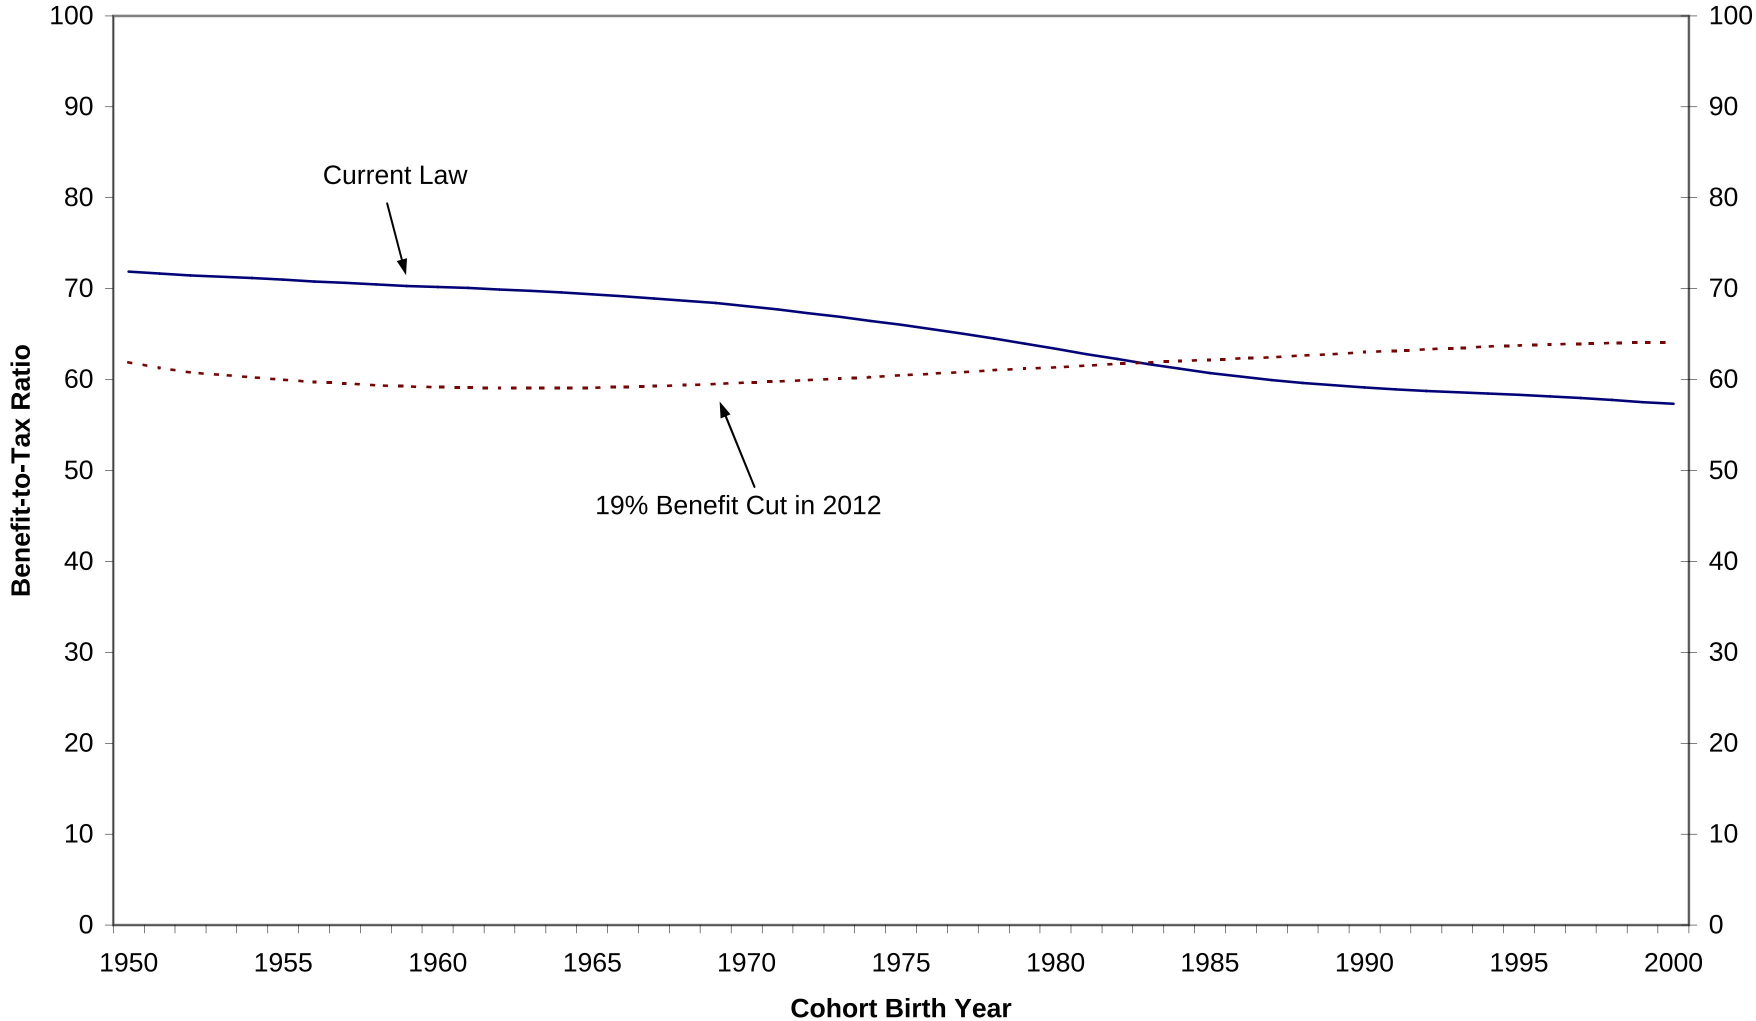

Figure 7. Potential Ranges for Ratios of Lifetime Benefits to Lifetime Taxes (Current Law and One-Time Sustainably Solvent Benefit Cut)

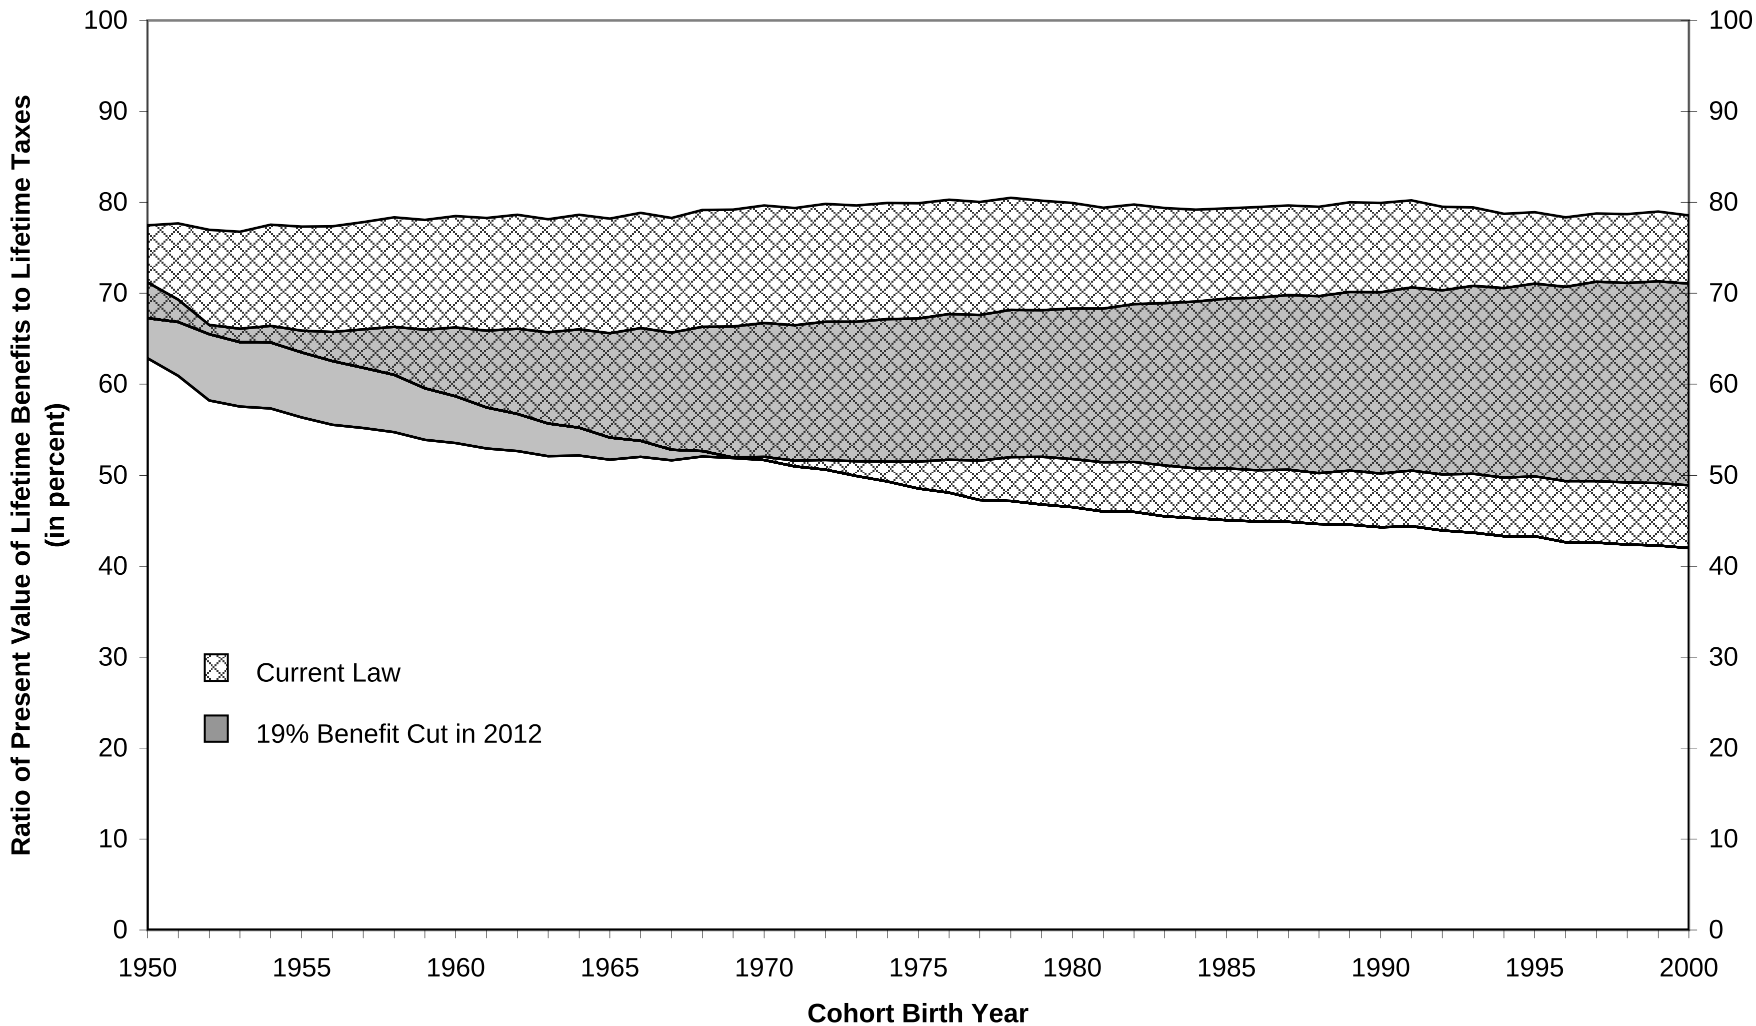

Figure 8. Potential Ranges for Social Security Trust Fund Ratios
(Current Law and One-Time Sustainably Solvent Benefit Cut)

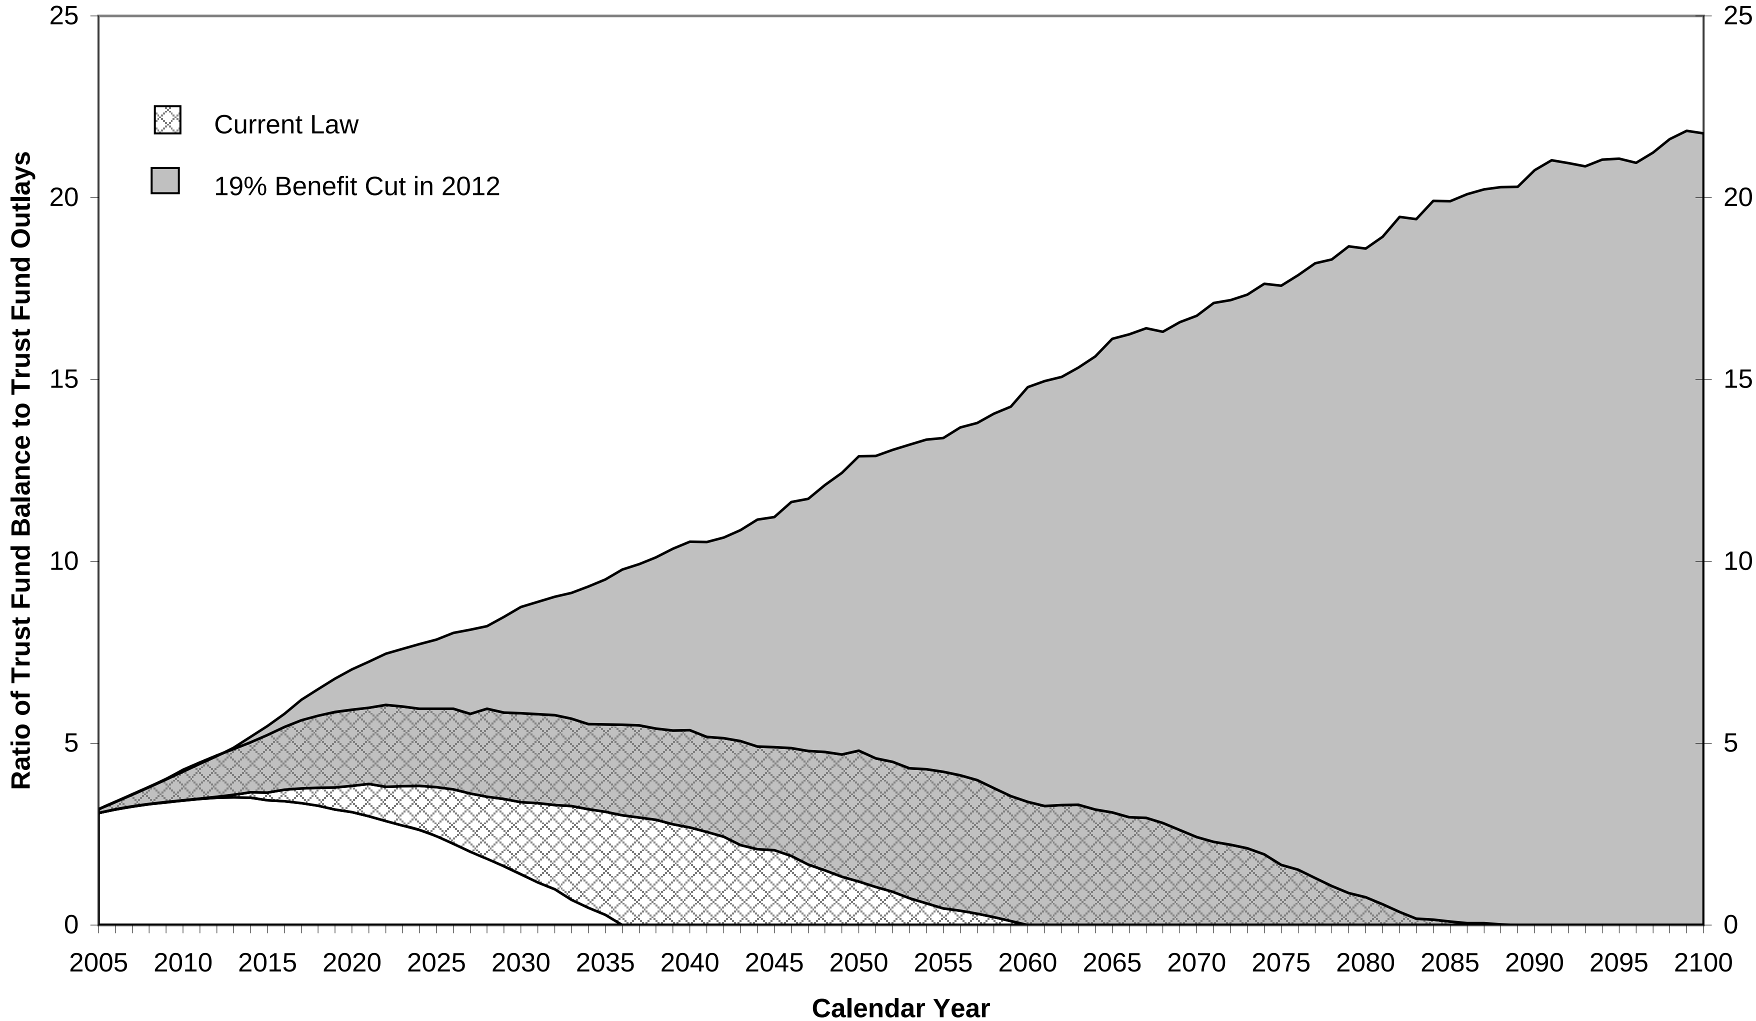

Figure 9. Potential Ranges for Ratio of Lifetime Benefits to Lifetime Taxes
(Sustainably Solvent Benefit Cuts with and without State-Dependent Policy Rules)

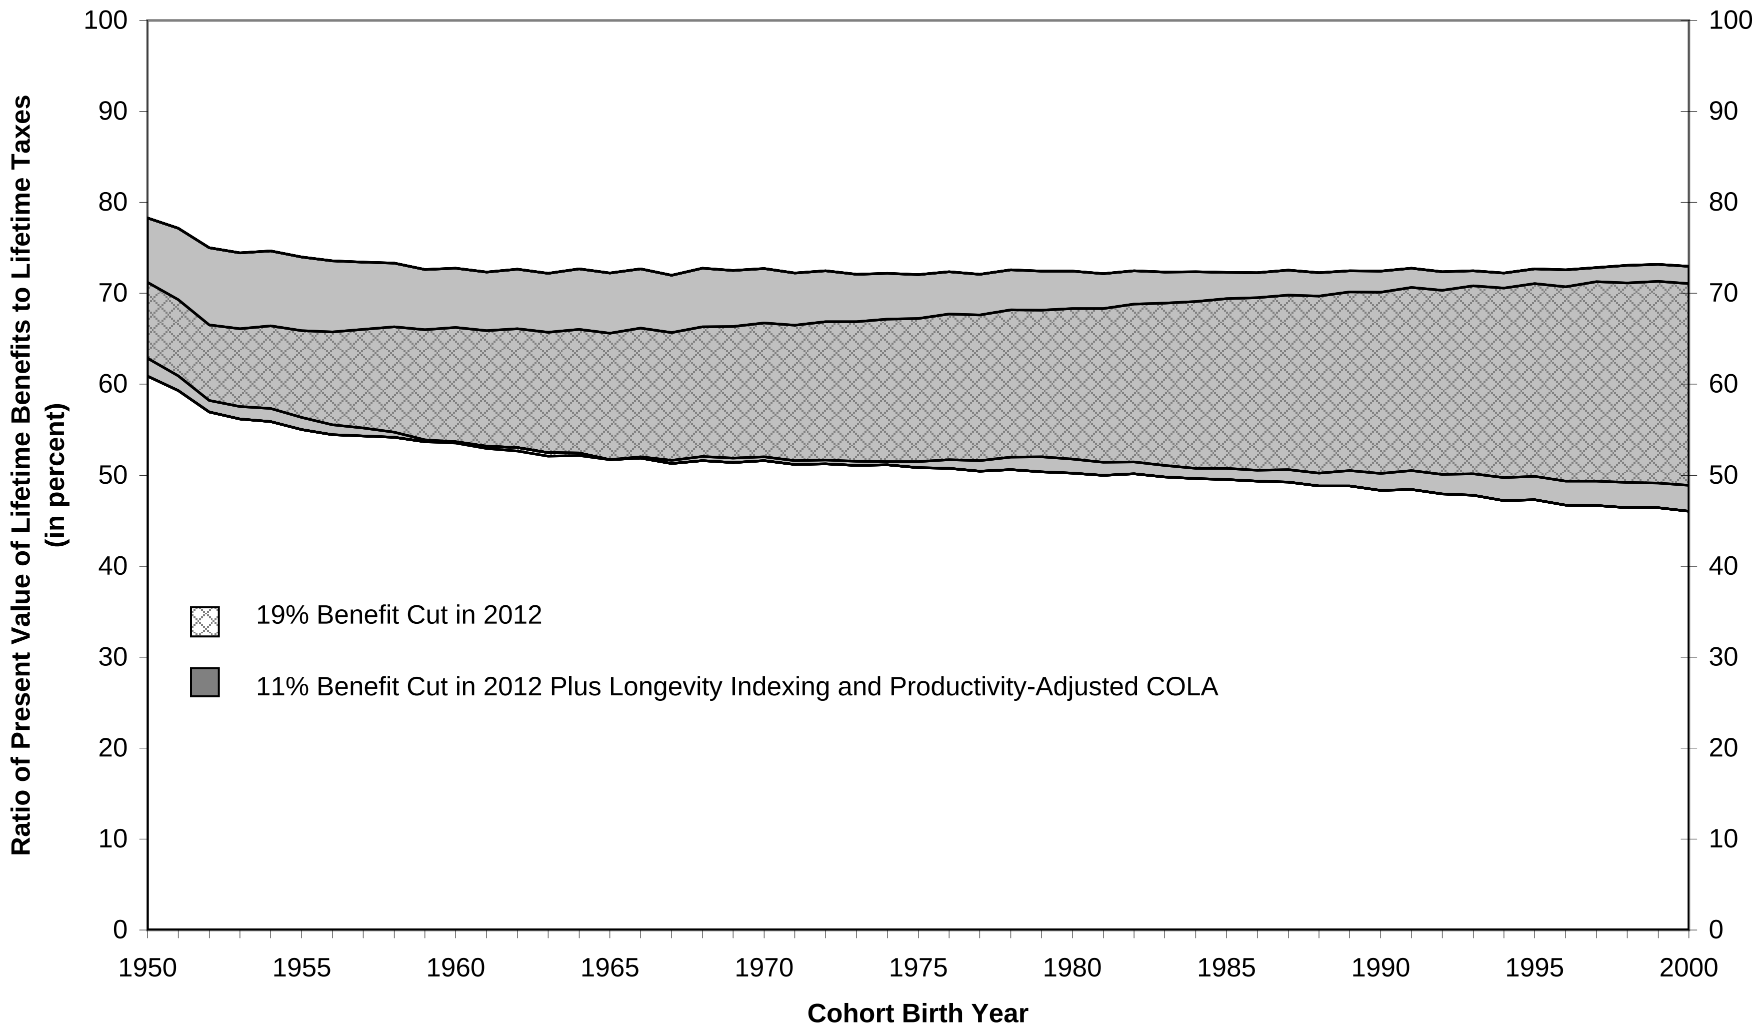

Figure 10. Potential Ranges for Social Security Trust Fund Ratios
(Sustainably Solvent Benefit Cuts with and without State-Dependent Policy Rules)

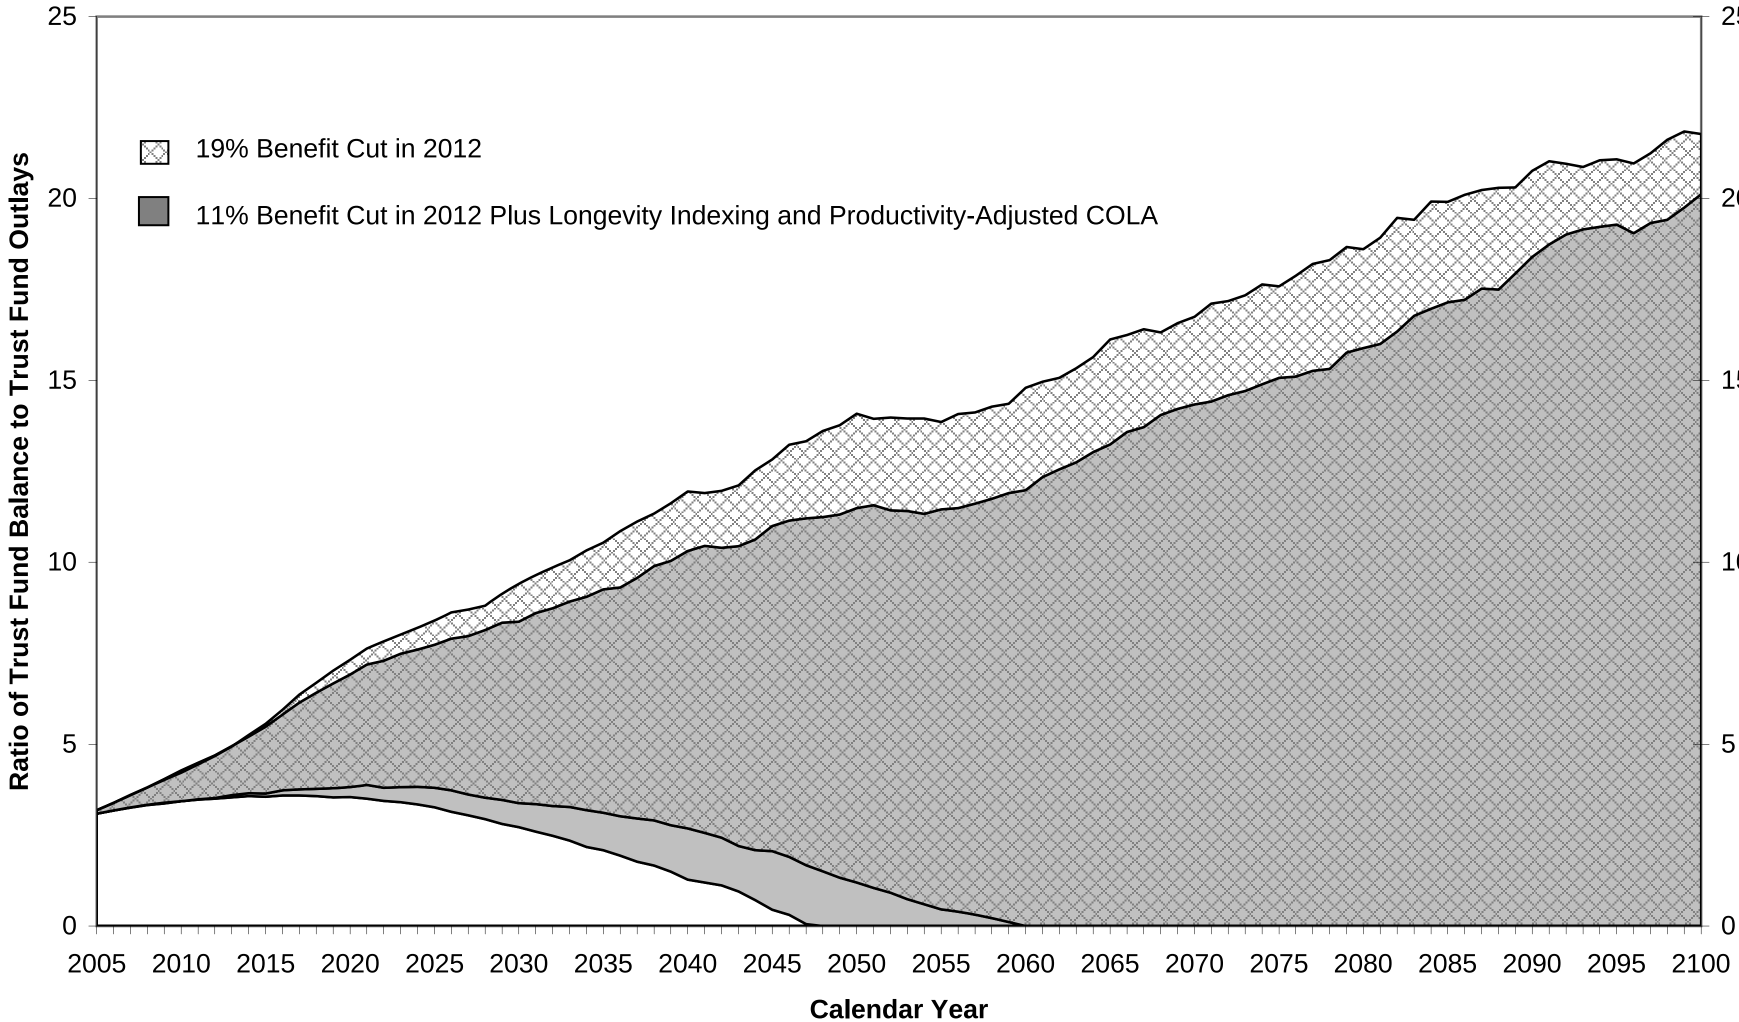

Figure 11. Potential Ranges for Ratios of Lifetime Benefits to Lifetime Taxes (State-Dependent Policy Rules with and without Fertility Variation)

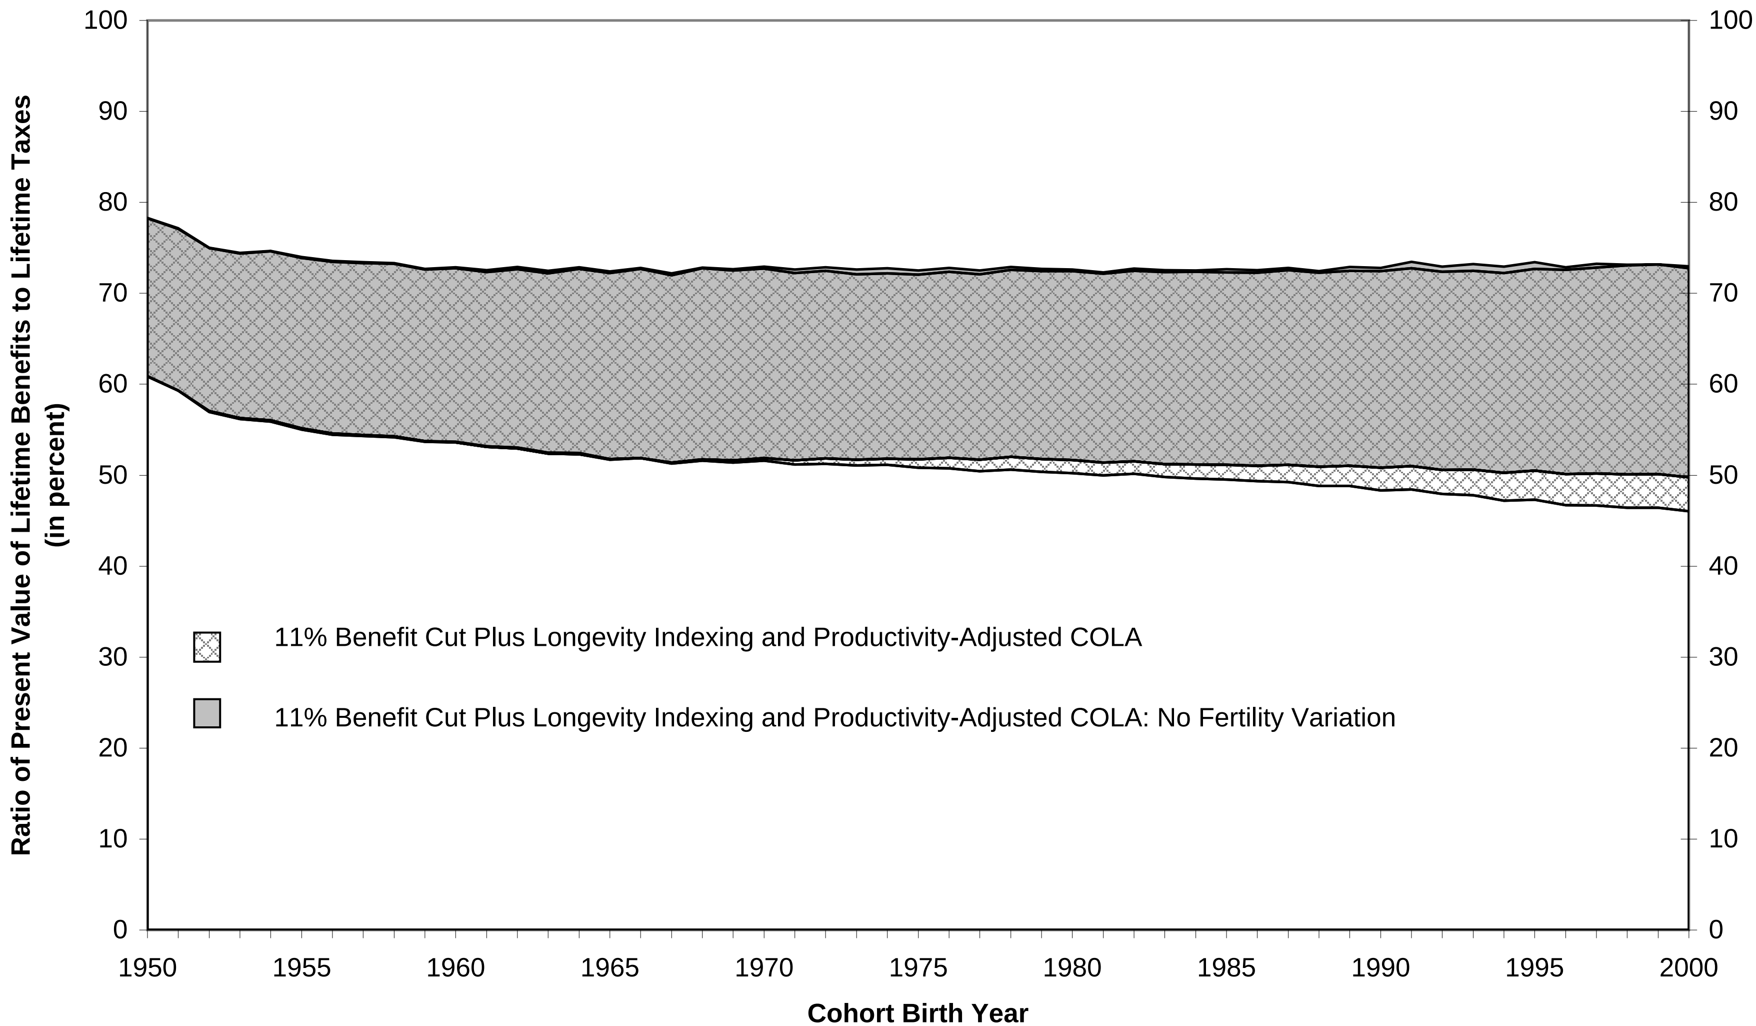

Figure 12. Potental Ranges for Social Security Trust Fund Ratios
(State-Dependent Policy Rules with and without Variation in Fertility)