The Market Value of Social Security

by

John Geanakoplos
Yale University

and

Stephen P. Zeldes
Graduate School of Business, Columbia University; and NBER

First draft: December 2007
Latest revision: September 29, 2008

The Market Value of Social Security

How much are the accrued benefits of workers worth today? How much is the Social Security system in deficit? The traditional actuarial answer is to ignore risk and compute expected value. If benefits are risky and this risk is "priced" by the market, then the actuarial estimates will differ from market value.

The exact adjustment for risk requires a careful examination of the stream of future benefits. The U.S. Social Security system is "wage-indexed", i.e. future benefits depend directly on the realization of the future economy-wide average wage index. We assume that there is a positive long-run correlation between average labor earnings and the stock market. We then use derivative pricing methods standard in the finance literature to compute the market price of individual claims on future benefits, which depend on age and on the macro state variables

We find that the difference between market valuation and "actuarial" valuation is large, especially when valuing the benefits of younger cohorts.

The market value of accrued benefits is only 2/3 of that implied by the actuarial approach. The market value of the 75-year deficit is 30% lower than the actuarial estimate.

John Geanakoplos James Tobin Professor of Economics Economics Department 30 Hillhouse Avenue Yale University New Haven, CT 06520 203-432-3397 john.geanakoplos@yale.edu cowles.econ.yale.edu/faculty/geanakoplos.htm	**Stephen P. Zeldes** Benjamin Rosen Professor of Economics and Finance Graduate School of Business Columbia University 3022 Broadway New York, NY 10027-6902 212-854-2492 stephen.zeldes@columbia.edu www.columbia.edu/~spz1

I. Introduction and related literature

The Office of the Chief Actuary in the Social Security Administration is regularly required to report on Social Security's financial status. The actuaries compute several measures of unfunded obligations, each of which requires computing the present value of streams of future benefits and (in some cases) contributions. These measures include the total cost of immediately transitioning to a personal account system ("maximum transition cost") and the difference between the present value of expected cash inflows and outflows ("the actuarial imbalance" or "open-group unfunded obligations"). These SSA estimates heavily influence policy makers' and the public's perception of the health of the Social Security system and also the likelihood and direction of reform.

While the methodology used by the SSA for arriving at these estimates is in many ways sophisticated, we argue that in one important way it is quite naïve: the methodology does not incorporate any adjustment for the market price of risk. When participants in financial markets value cash flows associated with traded financial assets, they adjust present values to incorporate the riskiness of the cash flows. All else equal, future cash inflows with greater risk (appropriately defined) are discounted with a higher discount rate, leading to a lower present value.

In this paper, we estimate the *market value* of Social Security's future expenditures (worker benefits) and receipts (worker contributions obtained from payroll taxes), i.e. the prices these future cash flows would command if they were traded as assets and liabilities in financial markets. We then construct market-based (i.e. risk-adjusted) estimates of three common measures of Social Security's financial health, and compare these estimates to those obtained using the traditional SSA methods with no correction for the price of risk. We argue that market-based valuations of this sort would provide useful information to government officials, financial market participants, and households.

For example, suppose that the Social Security benefits promised some worker were always equal to the dividends of one share of a particular stock. It would be sensible to quote the value of these benefits at the market price of the stock. That would for example allow the worker to compare the size of his savings portfolio, which might hold shares of the same stock, and his social security portfolio of benefits.

Similarly for the Social Security system as a whole, if all the promised benefits together were identical to 20% of the combined European stock market, then the price of 1/5th of the European stock market would be a useful guide to understanding the cost of transitioning to a fully funded Social Security system. Under the current methodology, however, the SSA would likely report much larger numbers for the worker's promised benefits and for the system as a whole, because those numbers would ignore the riskiness of the dividends. We would like to estimate precisely how much difference a market calculation makes. This last example is not so far-fetched as it might appear. Payroll taxes are proportional to wages, up to some maximum. And under current rules, Social Security benefits for a worker are tied to the economy-wide average wage at his retirement. (We discuss the precise formula later). In the long run, wages and dividends must move together. If it were clear that fifty years from now American businesses would be failing and paying no dividends, we could be pretty confident that wages would be low by then as well.

Historically the total return on stocks has been much higher than the riskless rate. This suggests that dividends are indeed subject to the kind of uncertainty that leads cash flows to be more heavily discounted by the market. Hence both Social Security expenditures and Social Security receipts would be valued less by the market than by the SSA methodology if they were traded.

The maximum transition cost is meant to measure the value of all the Social Security benefits accrued to date by workers living today. According to the SSA, that figure stands at $17.6 trillion as of early 2008. By applying the risk adjustment we shall discuss later, we calculate a market value about 25% lower.

When it comes to valuing the difference between benefits and tax receipts over the next 75 years (the open group unfunded liabilities) the situation is qualitatively more subtle. Which number should go down more as a result of the risk correction, the taxes or the benefits?

The answer is the taxes. As of any date t in the future, the taxes depend on the wages in the same period t. But the benefits of each worker at time t depend on the wages at his retirement at time s < t. Thus the wages are more correlated with dividends than are the benefits.

4

Nevertheless, though receipts are discounted more than benefits, both shrink, as does their difference. The SSA estimates the open group unfunded liabilities at $4.3 trillion. After the risk correction, we estimate a number about 30% lower.

In order to focus on the key conceptual points, we use a relatively simple stochastic model for our estimates. We assume that there is only one source of uncertainty in the cash flows: average economy-wide earnings, or what SSA calls the "Average Wage Index" (AWI). This uncertainty is important because by law future benefits depend explicitly on the AWI at the time of retirement, and future contributions in each working year are proportional to earnings in the corresponding year. We assume processes for earnings and stock prices such that the AWI is highly correlated with the stock market over long horizons, but uncorrelated over short horizons. We then use modern financial tools (risk neutral pricing) to price Social Security cash flows as derivatives on the stock market. Along the way, we use the framework from our previous work on Progressive Personal Accounts (Geanakoplos and Zeldes, 2008) to describe accrued benefits in terms of units of a potentially tradable financial security that we called a PAAW (for Personal Annuitized Average Wage security).

We find important differences between our market-based estimates and those constructed using traditional methodology. As we said, our "maximum transition cost" estimate is about 25 percent smaller than the expected present values calculated using SSA methodology, and a market-based estimate of the "open group unfunded liability" is about 30 percent smaller than the SSA figure. The market-based Closed Group Unfunded Liability is *less than half* the size of the official figure.

Our paper is structured as follows. Section II describes the Social Security rules on taxes and benefits, and defines three measures commonly used to gauge the financial status of social security: "maximum transition cost", "closed group transition cost", and "open group unfunded obligations." In section III we present our arguments in favor of using market-based measures to evaluate Social Security's financial status and respond to arguments against this approach. Section IV presents our model and describes our methodology for computing market value estimates of these measures. In Section V, we present results for the three measures of Social Security's financial status. Section VI concludes and describes paths for future work. In particular, we

argue that our approach could be extended to more complicated and sophisticated stochastic models of the macro-economy and asset prices.

II. Program rules and actuarial measures of financial status

A. Social Security program rules

Under current Social Security rules, workers and employers together contribute 12.4% of "covered earnings" (i.e. all labor income below the earnings cap, equal to $102,000 in 2008). Of the 12.4 percent, 1.8 percentage points are earmarked for disability coverage. In the analysis that follows, we ignore DI coverage, and therefore use a social security contribution rate of 10.6 percent. Note that this implies that aggregate contributions in any year will be proportional to aggregate covered earnings.

Upon retirement, workers receive benefits that are linked to their history of covered earnings. Geanakoplos and Zeldes (2008) show that because the system is "wage indexed," it can be described more easily and clearly by defining a set of "relative" variables that are equal to the dollar amounts divided by average economy-wide earnings for the year. We define *relative earnings* for a worker in any year t as his current covered earnings for that year divided by average economy-wide earnings, and *average relative earnings* as the average of his highest 35 values of relative earnings. *Initial relative benefits* are defined by a concave function of the individual's average relative earnings. *A worker's initial dollar benefits* (also referred to as the Primary Insurance Amount or PIA) equals initial relative benefits multiplied by average economy-wide earnings (the "average wage index" or AWI) in the year the worker turns 60.[1] Benefits in subsequent years are indexed to the CPI, so that individuals receive a constant stream of real benefits for as long as they live. As we shall see, the fact that benefits are directly related to the average economy-wide wage index in the computation year leads directly to the risk adjustment that we perform when we compute market valuations.

B. Measures of financial status

We examine the three most common measures of the actuarial status of Social

---

[1] Workers who start collecting benefits at the Normal Retirement Age (65 for workers born before 1938, increasing to 67 for workers born after 1959) receive the PIA as their benefit. Benefits are lower for those who start collecting at a younger age and higher for those who start collecting at an older age.

Security: the Maximum Transition Cost, the Closed Group Liability, and the Open Group Liability.

The Maximum Transition Cost measures only benefits accrued to date, thus ignoring any future accruals. Calculated by taking the difference between benefits accrued to date and Trust Fund assets, the Maximum Transition Cost is the unfunded obligation measure that corresponds most closely to shortfall measures used for private pensions. As of 2008 this shortfall stands at $17.6 trillion. To obtain the Maximum Transition Cost implied by the model, we need a definition of "accrued benefits." By definition, accrued benefits can rise, but never fall (under the accrual rule). In Geanakoplos and Zeldes (2008), we show that there are a variety of feasible accrual rules and describe two natural rules in detail. For simplicity, we focus here on the "straight-line" method that corresponds closely to that used by Wade, Schultz, and Goss (1998).

The Closed Group Unfunded Liability is a 75-year measure of the difference between future benefits that will be paid to individuals that are already working and the taxes that will be paid by these workers. The closed group liability differs from the Maximum Transition Cost because it includes the future tax payments and accruals of current Social Security participants. Workers are assumed to contribute and draw benefits according to the current formulas for the remainder of their lives. By this measure, the program's unfunded liability over the next 75 years is $15.2 trillion.

The Open Group Unfunded Liability considers all prospective flows over the next 75 years, including taxes and benefits for those not yet in the workforce. This number is significantly smaller than the closed group figure because it includes new workers who enter the system late enough that their tax payments fall within the 75-year forecast horizon, but the benefits they will receive do not. The open group liability currently stands at $4.3 trillion.

For more details on these measures, see Goss (1999) and Wade, Schultz, and Goss (1998).

III. The importance of market valuation

There has been considerable debate among actuaries, accountants, and economists about the merits of actuarial value versus market or "fair" value accounting

measures. A market price for social security liabilities would provide important information to households, governments, private pension plans, and other market participants. For government administrators of Social Security, the ability to determine the market value of the social security liabilities also implies the ability to hedge them (since valuation and hedging are dual computations). If the Social Security trust fund were someday permitted to diversify out of government bonds, this would provide a valuable guide to determining the optimal portfolio allocation. Indeed, the existence of a hedging opportunity would encourage such a law change.

For individuals, a market price for the benefits accrued by their cohort would provide information about the market value of their own benefits, helping them with financial planning decisions regarding saving and asset allocation. A market value for benefits would also make it more difficult for the government to take them away, thus further enhancing property rights.

Additionally, the benefits to an individual worker become closer to a (real) annuity as she nears retirement. Our valuation methods could be adapted to the pricing of individual annuities and of securities tied to aggregate longevity.

Finally, as we have argued elsewhere (Geanakoplos and Zeldes, in progress), it is conceivable that Social Security benefits will be traded in the future. Buyers and sellers of these new securities would be forced to make the same kind of computations we propose here. The government conceivably could purchase securities from the private sector that would replicate some of the benefits promised by the social security system.

FASAB

The Federal Accounting Standards Advisory Board (FASAB) has published a preliminary statement on new standards for Social Insurance Accounting for entries on the Balance Sheet of the United States (FASAB, 2006). The document describes two views. The Primary View, held by the majority of the board, would recognize every accrued benefit as a liability of the system.[2] Under this view, liabilities should be based on expected benefits "attributable" to earnings to date, using current benefit formulas. In contrast, the Alternative View advocates continuing the current practice of

---

[2] Accrued benefits would be those earned by fully-insured participants (e.g. social security participants who have achieved 40-quarters of covered earnings, the minimum to receive benefits) based on their earnings histories to date.

8

acknowledging only those benefits that are "due and payable" at time of valuation. Essentially, under the alternative view only current-period benefits not yet paid to beneficiaries (an amount close to zero) would be counted as a liability.

Supporters of the Primary View argue that recognizing the new liability is most consistent with the principal of accounting based on accrual, as opposed to cash flows, and best captures the economic costs incurred by social insurance programs each year. Supporters of the alternative view argue that given political and economic uncertainty regarding Social Security, such obligations are neither legally guaranteed nor reliably estimable. They also worry that, because of the large size of the obligation, incorporating it as a liability may make other important spending choices appear inconsequential.

Whether or not one wishes to characterize future benefit obligations as "liabilities", it is useful to compute their present value. Even the Alternative View would continue to require the Social Security Administration to report the present value of future cash flows, although not on the balance sheet. If the Primary View is adopted, the results of our paper will be particularly important, because the Social Security Administration and Office of the Actuary will be charged with the task of computing the present value of accrued benefits and reporting it on the balance sheet. Proper valuation of these risky liabilities will be essential to the new guidelines' efficacy in accurately portraying the financial status of the Social Security program. Additionally, the fact that risk-corrected valuations are significantly smaller than standard risk-free valuations may reduce (although not eliminate) the concern that the size of the accrued benefit obligation will overwhelm other information in the balance sheet.

IV. Model

In this section, we describe our model for dividends and wages and how we price wage bonds and PAAWs. We then describe our demographic model and how it is used to compute actuarial measures of the financial status of social security.

A. Dividends and wages

Social security taxes and benefits are tied to the path of future average economy-wide wages. Lucas and Zeldes (2006) showed how to value defined-benefit pension

liabilities when payouts are tied to future wages of the individual. We apply a similar approach here, modifying it to take into account the specifics of the Social Security benefit and tax rules.[3]

To proceed, we need a macroeconomic model that links wages and dividends. We would like our model to take into account that the correlation between wage growth and stock returns is low or zero at short horizons, but higher at longer horizons. We adopt the model used in Benzoni, et al, (2007), which generates a strong positive long-run relationship between average economy-wide wages and dividends (and therefore stock prices), but a short run correlation that is close to zero.

We suppose that dividends follow a geometric random walk. Using the discrete time analog of their continuous time process, we begin with a stationary random walk process for log dividends (d):

$$d_{t+h} - d_t = h(g_d - \frac{\sigma_d^2}{2}) + \sigma_d \sqrt{h} z_d \qquad (0.1)$$

The dividend growth shock, z_d , is assumed to be standard normal.[4]

Benzoni, et al assume a stationary pricing kernel with a constant price of risk, λ . This implies a constant price-dividend ratio, and therefore a constant dividend yield, $\delta = r + \lambda \sigma_d - g_d$. Because the stock price is proportional to current period dividends, it too will follow a geometric random walk with the growth in the stock price exactly equal to the growth in dividends. The total stock return thus equals the dividend rate *plus* the growth in dividends.

$$s_{t+h} - s_t \equiv (p_{t+h} - p_t) + (d_{t+h} - d_t) = h(g_d + \delta - \frac{\sigma_d^2}{2}) + \sigma_d \sqrt{h} z_d \qquad (0.2)$$

Note that (1.2) implies the counter-factual result that stock returns and dividend growth have the same volatility.

Next, consider the relationship between wages and dividends. Following Benzoni et al, we assume that wages and dividends are cointegrated. We write wage growth as function of 1) a deterministic wage growth, or "drift", parameter, 2) the

---

[3] One important difference is that under private DB pensions, the accrued benefit obligation (ABO) depends only on past labor earnings, while the promised benefit obligation (PBO) depends on future labor earnings. Due to the wage-indexing of Social Security, even the ABO measure depends on future (economy-wide) labor earnings.
[4] Equation (1.1) therefore implies a level representation with log-normal shocks and expected growth in

current-period deviation from the long term average wage-dividend ratio and 3) an i.i.d wage growth shock.

$$w_{t+h} - w_t = h(g_w - \frac{\sigma_g^2}{2}) - h\kappa(w_t - d_t - \overline{wd}) + v_w\sqrt{h}z_w \qquad (0.3)$$

Here, κ is an error-correction parameter, influencing how quickly wages tend to return to their long-run relationship with dividends. In general, we allow g_w to be different than g_d, although we will assume they are equal when we calibrate the model.[5]

Calibration

To calibrate the model, we need estimates of seven parameters: r_f, g_d, g_w, δ, κ, $\sigma^2{}_w$, and $\sigma^2{}_d$.

Where possible, we choose parameters that are consistent with the 2008 Trustee's Report intermediate cost assumptions. Accordingly, the risk free-rate, r, is set to 2.9% and average wage growth, g_w, to 1.1%. We further assume that wage growth and dividend growth are equal, implying g_d equal to 1.1% as well.

The expression for the dividend yield can be rearranged to give the equity premium as a function of the key level parameters, ($r_s - r_f = \lambda\sigma_d = \delta + g_d - r$). For our benchmark calibration, we choose the dividend yield in order to match the empirical equity premium. To estimate the equity premium, we observed that the average excess returns exhibited by the S&P 500 over 3-moth T-bills from 1959 through the first half of 2008 was 5.1%. This implies a counter-factually large dividend yield of 6.9% = 5.1% - 1.1% + 2.9%. In the future we will recalibrate our results using a lower risk free rate and a higher dividend growth rate. That will give a lower dividend yield.

As noted by Benzoni et al, the slow reversion of wages to their long-run relationship means κ is difficult to estimate. As a baseline, we follow those authors by setting κ equal to .15.

Finally, we must parameterize the volatilities of dividends and wages. Because our model implies that dividend growth shocks carry through one-for-one to stock

---

the level of dividends equal to g_d.

[5] Benzoni et al derive a similar process by directly positing a process for the deviation from steady-state and then deriving the implied wage process. We generalize their model by allowing $g_d \neq g_w$ and restrict it by assuming the dividend shock, z_1, is orthogonal to the wage process. Otherwise, these two approaches are equivalent.

returns, we let σ_d equal 12%, the return volatility of the S&P over the period since 1959. Finally, we let σ_w equal 2%, the volatility of hourly manufacturing wage growth over the same period.

With these parameters there is a unique set of stationary risk neutral probabilities that maintain the same log normal volatilities but also rule out arbitrage between the stock and riskless bond. These probabilities will tend to weight paths where dividends decline more heavily than the objective probabilities do. Since in the long run wages are drawn to dividends, paths with declining wages will also get more weight from the risk neutral probabilities.

B. Pricing wage bonds

Following Lucas and Zeldes (2006), we assume that the risk unrelated to the stock market is unpriced, ignore aggregate longevity and interest rate risk, and use risk-neutral Monte Carlo derivative pricing techniques to price a set of wage bonds of different maturities as derivatives on the stock market. The wage bond prices will be a function of current average labor earnings and the current value of the stock market. We then use these prices to compute the market values of a set of PAAWs.

To generate wage bond prices, we generate 50,000 dividend and wage time series based on the risk neutral probabilities and processes and calibration described above. We then compute the price of a year-t wage bond by averaging across the simulated year-t wages, and discounting to the present using the risk free rate. Figure 1 shows the wage bond price for workers retiring over the next 75 years. For cohorts that retire in the next few years, the risk adjustment has little effect; their wages at retirement are already essentially fixed. However, for younger cohorts, the wage bond becomes increasingly risky, and the risk adjustment takes on high importance. For cohorts that will retire long into the future, present value discounting comes to dominate these effects, and both prices asymptote towards zero. Figure 2 shows that the ratio of the risk-adjusted to the unadjusted bond prices also goes to zero for bonds dated at long horizons.

C. Pricing PAAWs

In Geanakoplos and Zeldes (2008), we described how to create a system of

personal accounts that achieves many of the core goals of supporters of the current system, including redistribution and intergenerational risk-sharing. We called these "Progressive Personal Accounts." One step in that process was to show that personal accounts could be structured to exactly mimic the benefits promised under the current system. This involved the creation of a new financial security which we named a Personal Annuitized Average Wage security, or PAAW for short. Whether or not Progressive Personal Accounts are adopted, we can use the (as yet, theoretical) PAAW security to describe and value Social Security obligations.

We define a Personal Annuitized Average Wage security or PAAW as a security that pays its owner one inflation-corrected dollar for every year of his life after a fixed date t_R (the year he hits the statutory retirement age R), multiplied by the economy-wide average wage at t_R. PAAWs are tied to specific individuals (i) and to the year of the first payout on the security (t_R), and we use the notation PAAW(i,t_R) to capture this.

Each additional dollar that an individual earns generates additional accrued benefits or PAAWs. At any point in time t, an individual's accrued benefits can be summarized completely by the number of PAAWs owned. The present value of accrued benefits is therefore equal to the quantity of accrued PAAWs (known at time t) multiplied by the price of a PAAW(i, t_R). Here, we describe how we compute the price of PAAWs, and in the next subsection we estimate the quantity of PAAWs outstanding for each cohort.

Since per-period benefit payment are fixed (conditional on survival) at retirement, the price of a PAAW is simply the price of a wage bond times the value of a $1 lifetime annuity starting in the retirement year. We therefore generate the aggregate PAAW(t) price by discounting such an annuity using the risk-free rate and cohort-specific survival probabilities through age 119, and multiplying the result by the time-t wage bond price. Our mortality assumptions are based on the cohort life tables from Social Security Actuarial Study 116. We assume for now that all individuals of the same age face the same conditional survival probabilities[6], i.e. that there is no heterogeneity or private information about these probabilities.

Figure 3 shows per-PAAW prices for workers from different current-age cohorts. Since the value of a PAAAW(t) is just the present value of an annuity stream equal to

---

[6] For the calculations presented, we used the survival probabilities for males born in 1980.

the time-t wage bond, the patterns in these figures are tightly linked. The unadjusted profile is increasing for younger workers because the risk-free discount rate exceeds average wage growth; the risk-corrected profile is even more steep because of reduced riskiness in wages as workers approach retirement. After retirement, the profile become downward sloping because of both reduced life expectancy and the fact these benefits were earned at lower wages on average.

Figure 4 shows the ratio of risk-corrected to naïve (non-risk adjusted) PAAW prices for each cohort. For cohorts that have already retired, the risk-adjustment has no impact on the valuation. For younger cohorts, however, there is a significant difference between the two methods. For cohort aged 41 and younger, the risk-adjusted measure is less than half of the naïve valuation.

D. Future taxes and benefit flows

In order to compute estimates of the actuarial balance other than the transition cost, it is necessary to estimate future Social Security benefit payments and tax income. To do this, we combine our Monte Carlo model for aggregate wages with simple models for demographics and individual PAAWs accrual. The construction of this model is detailed in the appendix. Figures 5 and 6 show that with simple assumptions about population growth and mortality we are able to get a close approximation to the population dynamics projected in the 2004 Social Security Trustee's Report. Our model does not account for immigration, which helps explain why – despite the fact that we closely match dependency ratios – we have a somewhat lower population growth rate in our model.

.

V. Results

A. Demographics and cash flows

We combine the series for benefits and taxes for each cohort in our Monte Carlo model to generate cohort-specific per-capita cash flow patterns for both the risk-adjusted and unadjusted scenarios**.** Figure 9 compares projected cash flow measures for the simulated economy to projections from the Social Security Trustee's report. Because the scale of the simulation is arbitrary, we scale these cash flows to match the magnitude of the SSA report in the first year. Given the simplicity of our model and our

14

demographic assumptions, these series match remarkably well. Because the relative magnitudes and timing of our flows match well, our model should give reliable results regarding the important of using a risk-adjustment in pricing the liability.

B. Maximum transition cost

Figure 7 graphs the quantity of outstanding PAAWs as function of age. No simulation is needed to compute these quantities. Since our model does not account for spouses and others earning benefits based on other earnings histories, we normalize initial population to 150 million – based on a rough estimate that approximately half of US citizens will work a full career under Social Security. This normalization is relevant for Figures 7 and 8. Estimating the market value of accrued benefits simply requires multiplying these quantities with prices from our monte carlo pricing simulation, and summing across cohorts. Summing up the values in Figure 8 gives an accrued liability of 9.9 trillion, which is somewhat less than the 2005 Trustee's Report estimate (adjusted to excluded spouses, survivors and disability) of around 11.2 trillion. To match the maximum transition cost, we would need to subtract the value of the trust fund from both numbers. Since we cannot split out the portion of the trust fund that is related to the beneficiaries we model, however, this adjustment is not easy. In general, comparison of the absolute dollar amounts from the simulation and the SSA estimates are merely suggestive, and should not be taken literally.

More interesting, then, are the percentage adjustments suggested by the modeled economy. Table 1 shows that, overall, our measure of the accrued liabilities is around 3/4th of the unadjusted value. Figure 8 compares the risk-adjusted and the naïve (risk-neutral) valuations by cohort, using the OACT-based estimates of quantity, respectively. Note that the risk-adjustment reduces the value of the liability for all of non-retired cohorts, and this includes the five cohorts with the largest number of PAAWs accumulated (ages 40 – 60). Thus, risk correction significantly changes our quantitative assessment of the value of accrued Social Security benefits.

C. Closed group

The second row of Table 1 shows our measures for both the closed and open group liabilities, using our simple model of population dynamics and PAAWs accruals.

The risk adjustment dramatically reduces our estimate of the closed group liability, from about 2.8 to 1.4 times payroll, an adjustment of 50%. The adjustment is even more important for the closed group liability than it is for the accrued liability measures.

Why is this effect so strong? To see why, consider any currently-working cohort (recall that benefits are fixed for retired workers regardless). From the perspective of that cohort, earnings always come prior to benefits; therefore cohort earnings are less risky. Because benefits for each cohort are relatively more risky, they will be disproportionately impacted by the risk adjustment – and the present value of the imbalance from the perspective of that cohort will be reduced. In the closed group, every non-retired cohort works in this direction. When the initial (unadjusted) estimate is negative, this effect is reinforced by the fact the risk-adjustment shrinks the value of any given cash-flow toward zero.

D. Open group

Figure 10 compares the present value of tax and benefits generated by our model for each of the next 75 years, in both risk-adjusted and unadjusted terms. The graph shows the system starts in surplus, but quickly turns to deficit. Furthermore, the deficit shows up nearly ten years earlier under the risk-adjusted measures. Another key observation, which we return to later, is that the risk-adjust flows always smaller than their unadjusted counterparts, and tend toward zero much more quickly.

The 75-year open group liability (see Table 1), sums the present values in Figure 10. This liability is smaller than the closed group liability under both measures, and the risk-adjustment is less important in percentage terms. The key difference between the open and closed group measures is that the open-group includes a large number of workers who have earnings (and therefore contribute taxes) over the later years in the horizon, but don't collect benefits until the 76[th] year or later. Because these inflows are not matched by benefit payments, this decreases the future flow deficits (the fundamental nature of pay-as-you-go) and thus the computed liability.

Understanding why risk adjustment has the impact it does, requires distinguishing between two different effects of the adjustment. The first effect, we will call the "tax effect," stems from the fact that the taxes earned in the open group scenario are subject to a larger risk correction than are benefits. To see this, consider

16

the risk adjustment for any calendar year t. In that year, some participants will be beneficiaries; benefits must have been determined at t-1 or (for most benefits) even earlier. Other participants will be paying taxes, which depend on wages that are determined only in time t exactly. Wages are therefore more risky and subject to a larger risk correction. This effect always works to increase the measured liability in the system. This "year-by-year" viewpoint is essential to understanding the open group measure, while the "cohort" viewpoint was preferable for understanding the closed group measure.

The second effect stems from the fact that the risk correction reduces the magnitude of *both* taxes and benefits. The price of any asset or liability is reduced towards zero if it is risky. We call this the imbalance effect, because it increases in importance the larger the unadjusted imbalance, positive or negative. Therefore, unlike the tax effect, the direction of the imbalance effect depends on the starting point. For a system in surplus, the effect of risk adjustment tends to shrink the surplus, making the system look less well off. On the other hand, a system in deficit will look better off due the impact of the imbalance effect.

Figure 11 gives the net present value of future cash flows for the simulated Social Security system. Interestingly, the risk adjustment makes the imbalance in the early portion of simulation look *worse*. Because the simulation starts in a deficit, the tax effect is dominating the imbalance effect. In the later period, however, the situation is reversed; as the unadjusted deficit become larger the imbalance effect comes to dominate the tax effect, leading to lower measures of the deficit. The fact that our open group measure shrinks under the risk-adjustment tells that, overall, the imbalance effect is the dominant effect in our simulated economy. Figure 12 shows the ratio of the risk-adjusted to unadjusted cash flows that were presented in Figure 11.

Figure 13 shows how to decompose the difference between the adjusted and unadjusted measures concretely into these two competing (in this instance) forces. Conceptually, this is accomplished by decomposing the benefit series in the portion of benefits that are equal to taxes, and a residual. The lower line, denoted tax effect, shows the effect on the imbalance of adjusting taxes and only the portion of benefits equal to taxes – the hypothetical adjustment of a system in balance. The upper line, marked imbalance effect, shows the contribution of the adjustment to the "residual"

series, in other words, the adjustment caused by the imbalance in period flows.

E. The Importance of Demographics

An important question is how heavily our results depend on the particular population distribution and dynamics assumed in the simulation. For the (similar) Maximum Transition Cost and Closed Group measures, a stylized, steady-state version of the model (with no baby-boom) yields only a slightly smaller reduction in the closed group figure. The open group measure is somewhat more impacted by adjusting the population dynamics; the unadjusted values show a surplus in the stylized version of the model. In part, this is caused by the relatively small starting value for this adjustment, leading to large percentage swings. In general it makes sense that a baby-boom is not essential for the risk-adjustment to be important.

VI. Conclusions, policy implications, and future research

We argue that market value is the appropriate way to measure both assets and liabilities of the Social Security system. Market value calculations adjust correctly for risk, and differ in important ways from the standard actuarial approach that discounts expected cash flows with a risk-free rate and therefore does not adjust for risk. We estimate that adjusting for risk reduces the SSA measures of imbalance in the system between 25% and 50%.

There are at least three extensions that we could undertake in the very near future. First we could recalibrate our results in a model with a smaller equity premium, and smaller riskless interest rate. We chose the same riskless interest rate as the SSA, but that seems high by historical standards. Second, we could consider other sources of risk in wage volatility, besides dividend changes, that are also priced by the market. Third, we could try and match our model to data on spousal benefits. Fourth, we could consider longevity risk and demographic risk.

18

Appendix

In this appendix, we describe some additional details about our methodology.

A. Matching PAAW quantities- OACT Methodology continued

Wages projections are based on the 2007 Social Security Trustees Report "Intermediate Assumptions" and conditional survival probabilities are also based on the most recent OACT actuarial estimates. Our estimate of F is based on the Social Security Public Use (SSPU) dataset, a 1% sample of current retirees' benefits and income histories. To estimate F using the straight-line method, each year we calculate a benefit based on the worker's average earnings to-date and then multiply the resulting benefit by the fraction (age-18)/(62-18)[7]. In either case, individual accruals are then aggregated to get the percentage series, F. Implicitly, we are assuming the profile of PAAWs accrual is similar across cohorts.

The OACT's outgo figures include non-retirement beneficiaries, mainly spouses, who earn benefits based other participants' earning records. Since spouses and survivors receive a significant portion of benefits that do not appear in our other data source, we adjust our figures downwards so as to calculate only benefits being paid to workers based on their own earnings history. We first note that total payments (number of recipients * average benefit) made in 2001 to wives entitled based on age and to non-disabled widows aged 70-74 amounted to $3.8 and $7.8 billion, respectively (Social Security Administration 2002, Tables 5.F3 and 5.F11). We get total spousal payments by scaling up payments to wives to reflect the fact that payments to wives entitled based on age account for 96% of all payments paid to spouses of live workers (Table 5.F1). We therefore combine the payments to spouses and survivors to get total spousal and survivor benefits of $11.8 billion, or 15.3% of the 2001 total outlays to beneficiaries aged 70-74 in the OACT data. We reduce outlays to each cohort by 15.3% and use the adjusted results to calculate the number of PAAWs accrued by workers based on their own earnings history.

---

[7] The straight-line technique is closest to the calculation of the maximum transition cost (Goss 1999). Benefits are calculated in the same manner as disability benefits, for which the average of earnings to-date excludes one year of lowest earnings for every 5 years of earnings history (Compilation 2007). We allow workers to begin accruing benefits at age 18, ignoring the eligibility requirement of having a

B. Adjusting the Maximum Transition Cost

In the rest of the paper we will focus on valuing Social Security using the data described above. We will provide several comparisons of measures thus obtained with published measures of the Social Security Administration, including the maximum transition cost. Since our calculations focus on retirement benefits for individuals based on their own work history (i.e. excluding spousal benefits) we must adjust the 2004 maximum transition cost estimate of $13.5 trillion to see if our estimates are comparable. To make this adjustment we first add back the 2004 OASDI combined trust fund value of $1.7 trillion to get a total present value of all accrued benefits of $15.2 trillion (Goss et. al. 2008). Since this number includes disability benefits and payments to auxiliary beneficiaries like spouses and children we scale it down in order to compare it to our estimates of benefits accrued by workers based on their own earnings history. Multiplying by 68.84%, the fraction of all 2005 benefit payments made to retired workers based on their own earnings history (Social Security Administration 2006), we obtain an "official" estimate of the present value of accrued retirement benefits of $10.46 trillion.

C. Monte Carlo model

To get a simple but realistic model for future population dynamics we proceed as follows. First, we establish initial population weights for each birth year cohort, using US census data from 2005 on population by age. For each subsequent year, we generate new weights by 1) shrinking the population in each existing (already-born) cohort according the cohort-life tables in Social Security Actuarial Study #120, and 2) generating an initial population for the new cohort (new-borns) by assuming new births increase at a constant 1% rate.

To estimate PAAWs accruals for each individual we abstract away from earnings heterogeneity and use the average relative wage profile for men born in 1937. Since this profile exhibits some uninformative fluctuations from year to year, we fit profile shape using a quadratic function of time and use the smoother fitted values. The actual and fitted earnings profiles are shown in figure **X**.

Because workers are assumed to have identical earnings profiles, per-capita

---

minimum number of quarters of covered earnings.

period benefits (and therefore total benefit accruals) can be represented as the average wage at retirement times a single replacement rate, which is identical for all workers across all years. All workers begin working at age 20 and retire at age 65. Using this profile, and the straight line accrual method, we compute a replacement rate of around 48% of retirement year wages. Because benefits are a concave function of lifetime income, the average replacement rate in an economy with earning heterogeneity (but the same average profile) should be slightly higher.

Since the Social Security tax is a constant 10.6% of income, tax returns do depend on the distribution wages (with the exception of wages that exceed the maximum taxable income.) Per-capita tax payments are simply 10.6% of the average wage in each year.

Cash flows are discounted at the risk-free rate to obtain present values. The total value of cash flows for each cohort is simply the product of average cash flows and the population of the cohort.

D. Methodology for CWHS data

Our estimates from the CWHS sample, which is similar the SSPU except that it includes current *and* retired workers, are a cross-sectional version of our calculation for F. All of our PAAWs estimates are based on the Primary Insurance Amount (PIA) formula. For current workers, we estimate accrual by assuming no future earnings beyond 2004. This corresponds to the "fastest accrual method" defined in Geanakoplos and Zeldes, (2008). (Similar calculations could be done based on the "straight-line accrual method".) Note that both the OACT and CWHS methodologies include this assumption; it is incorporated into the OACT-based estimates via our estimate of F. For current retirees, PAAWs accrued are based on the final PIA, no future earnings assumptions are needed.

E. Comparing Estimates of PAAW quantities by cohort

We would not expect identical results from the two methods of estimating PAAWs accrual. The CWHS methodology is a more direct measure of our concept of accrual: individual wage histories are known and the PIA amount can be calculated with zero future earnings based on the standard formula. It can also be calculated on a year-by-year basis. The OACT data, on the other hand, gives only projected flows (form which accruals must be backed out) and is aggregated by 5-year cohorts. The

OACT data also contain payments to spouses and survivors that may not be perfectly excluded by the simple adjustment described above.

References

Benzoni, L, P. Collin Dufresne, and R. Goldstein, 2007, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets are Cointegrated," Journal of Finance, Vol 62, No 5, October.

"Compilation of the Social Security Laws, Including the Social Security Act, As Amended, Through January 1, 2007," Vol. 1,Title II, Section 215. Available at http://www.socialsecurity.gov/OP_Home/ssact/title02/0215.htm

Federal Accounting Standards Advisory Board, 2006, "Accounting for Social Insurance, Revised – Statement of Federal Financial Accounting Standards," available at http://www.fasab.gov/pdffiles/socialinsurance_pv.pdf, October

Geanakoplos, John, Olivia Mitchell, and Stephen P. Zeldes, 1999. "Social Security Money's Worth." in **Prospects for Social Security Reform**, edited by Olivia S. Mitchell, Robert J. Meyers, and Howard Young, Pension Research Council, University of Pennsylvania Press.

Geanakoplos, John, and Stephen P. Zeldes, 2008, "Reforming Social Security with Progressive Personal Accounts," NBER Working Paper # 13979. May.

Geanakoplos, John, and Stephen P. Zeldes, in progress, "Creating a Market to Trade Pension Liabilities."

Goss, Stephen C., 1999, "Measuring Solvency in the Social Security System". In Prospects for Social Security Reform, edited by Olivia S. Mitchell, Robert J. Myers, and Howard Young, Pension Research Council and University of Pennsylvania Press. pp. 16-36.

Lucas, Deborah and Stephen P. Zeldes, 2006, "Valuing and Hedging Defined Benefit Pension Obligations - The Role of Stocks Revisited," unpublished manuscript, Columbia University Graduate School of Business, September.

Social Security Administration, 2002. Annual Statistical Supplement to the Social Security Bulletin, 2002. Washington, DC: Office of Policy.

Social Security Administration, 2006. Annual Statistical Supplement to the Social Security Bulletin, 2006. Washington, DC: Office of Policy.

Wade, Alice, Jason Schultz, and Steve Goss, 2008. "Unfunded Obligation and Transition Cost for the OASDI Program," Actuarial Note #2008.1, Social Security Administration, Office of the Chief Actuary, July.
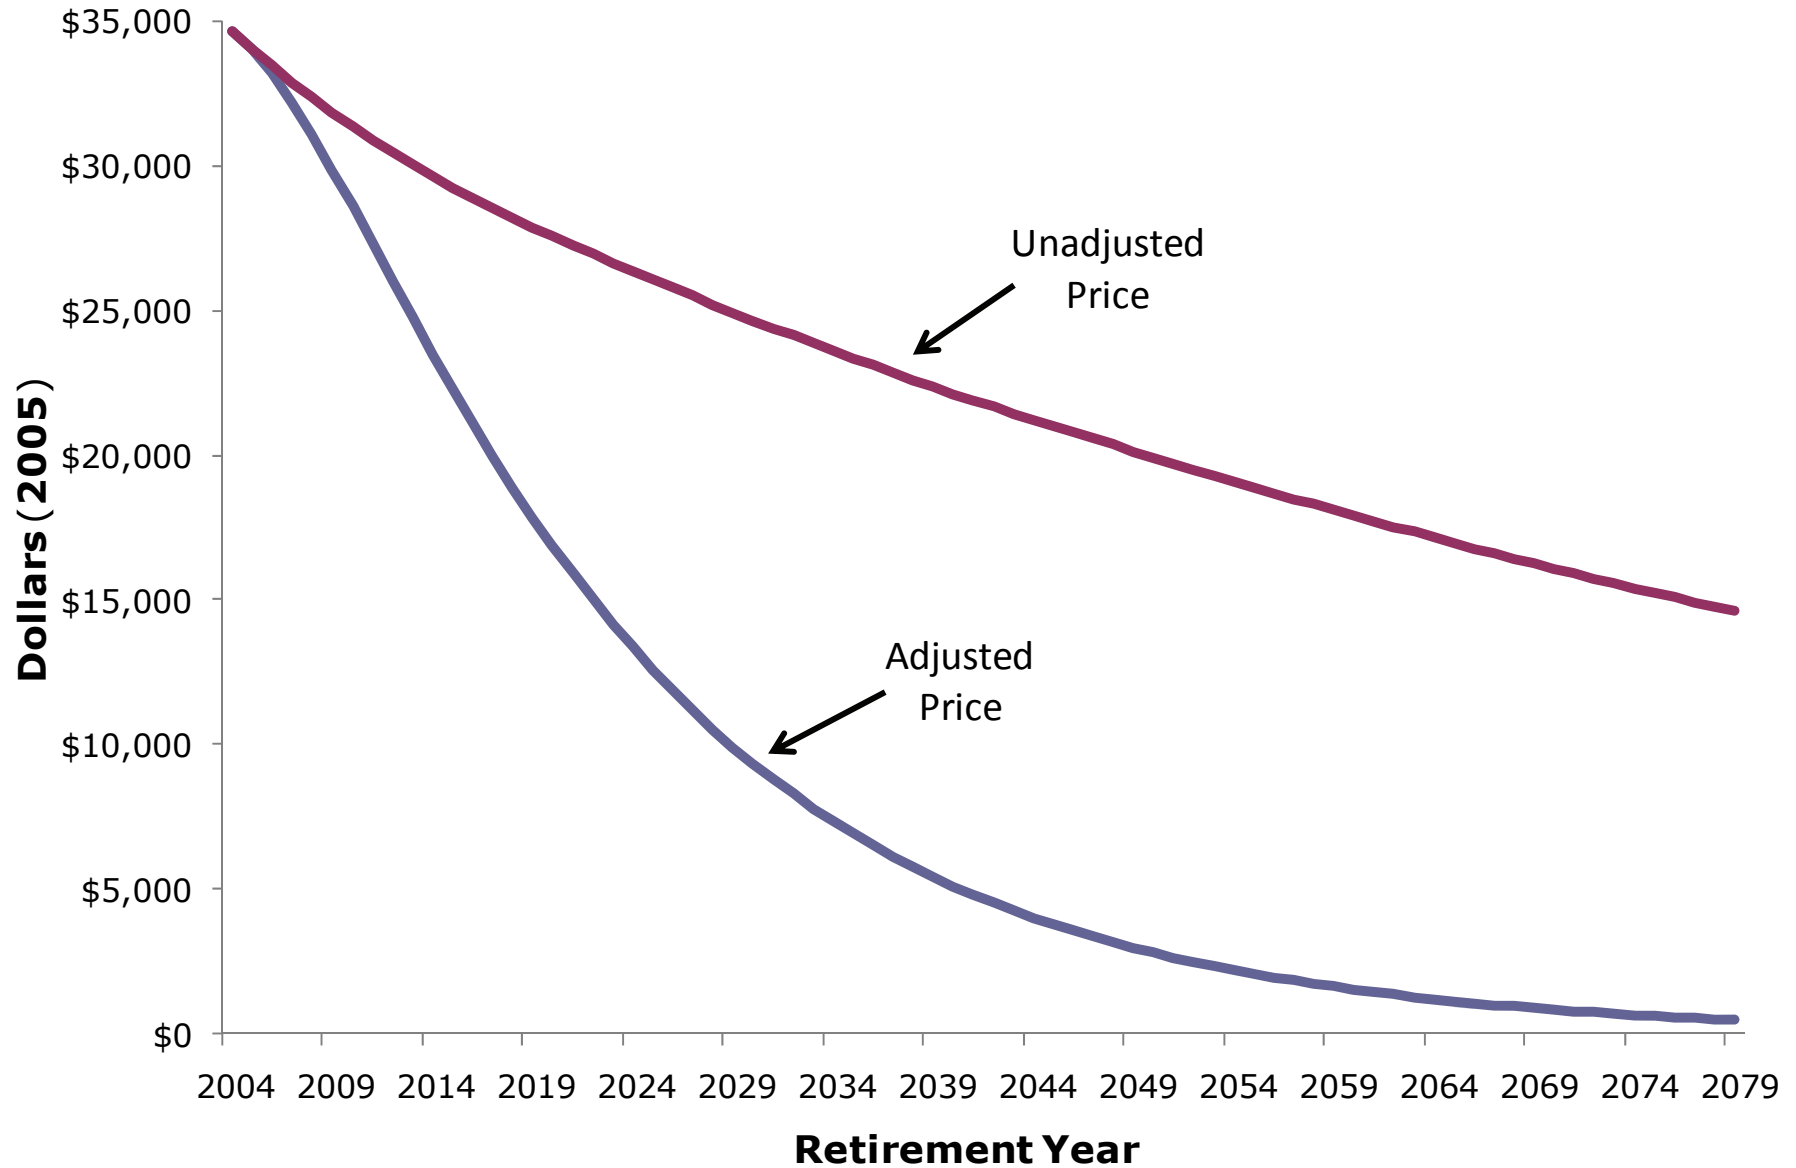
Figure 1: Wage Bond Price

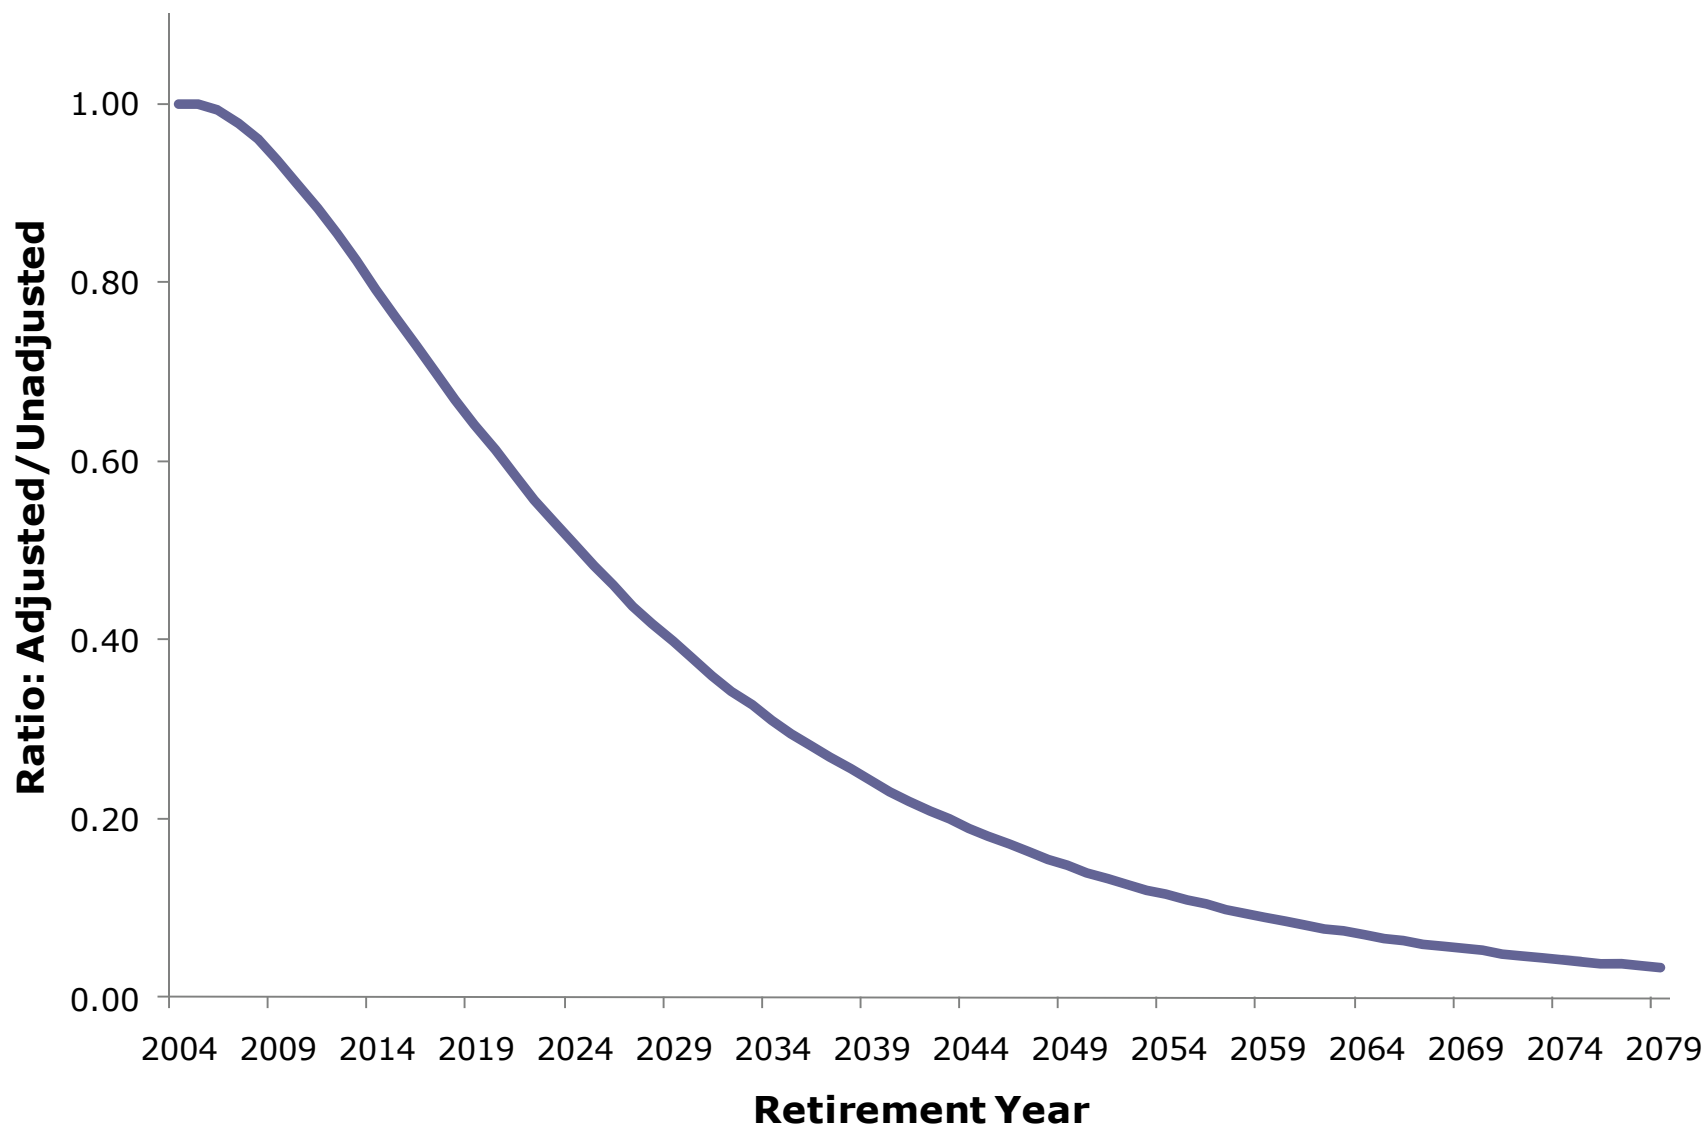**Figure 2: Ratio of Risk-Adjusted to Unadjusted Wage Bond Price**

Figure 3: Price per-PAAW

Naive
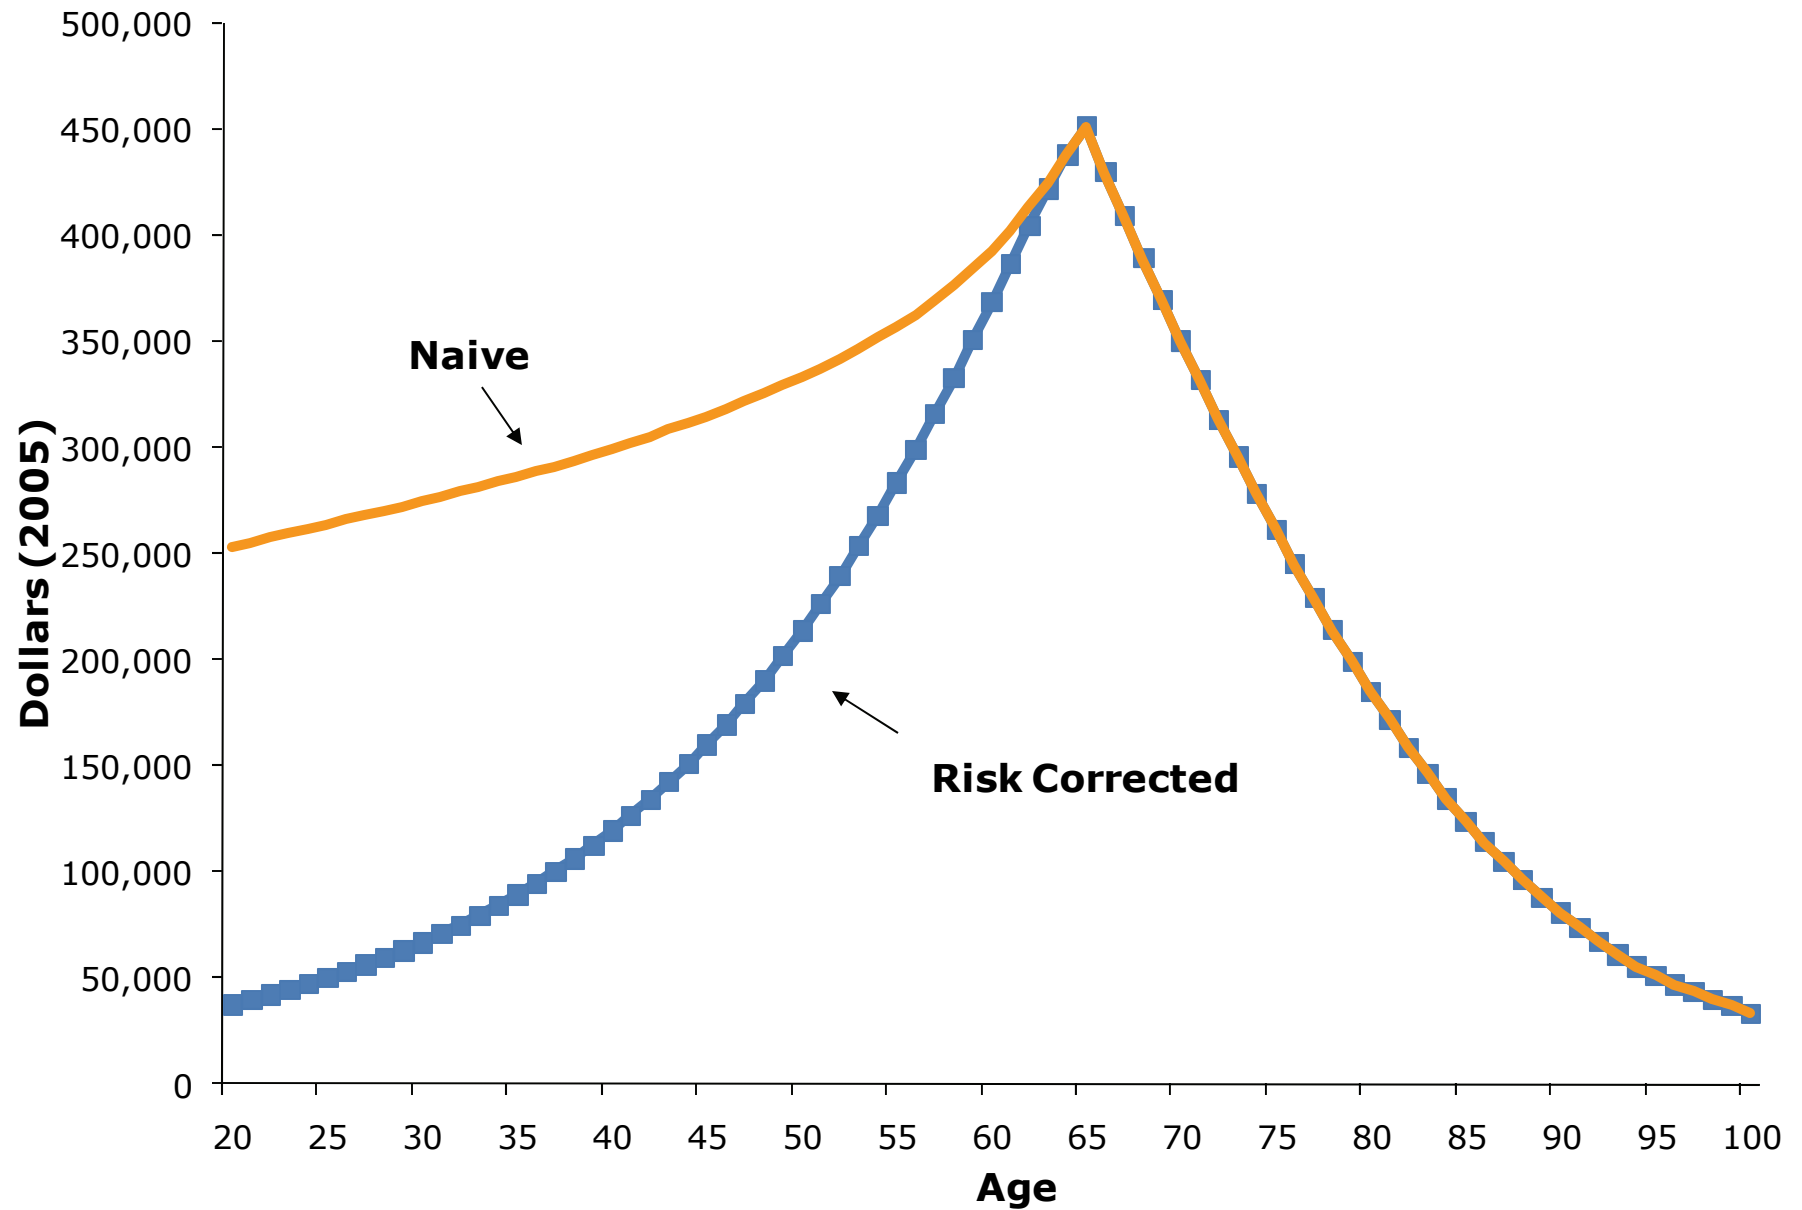
Risk Corrected

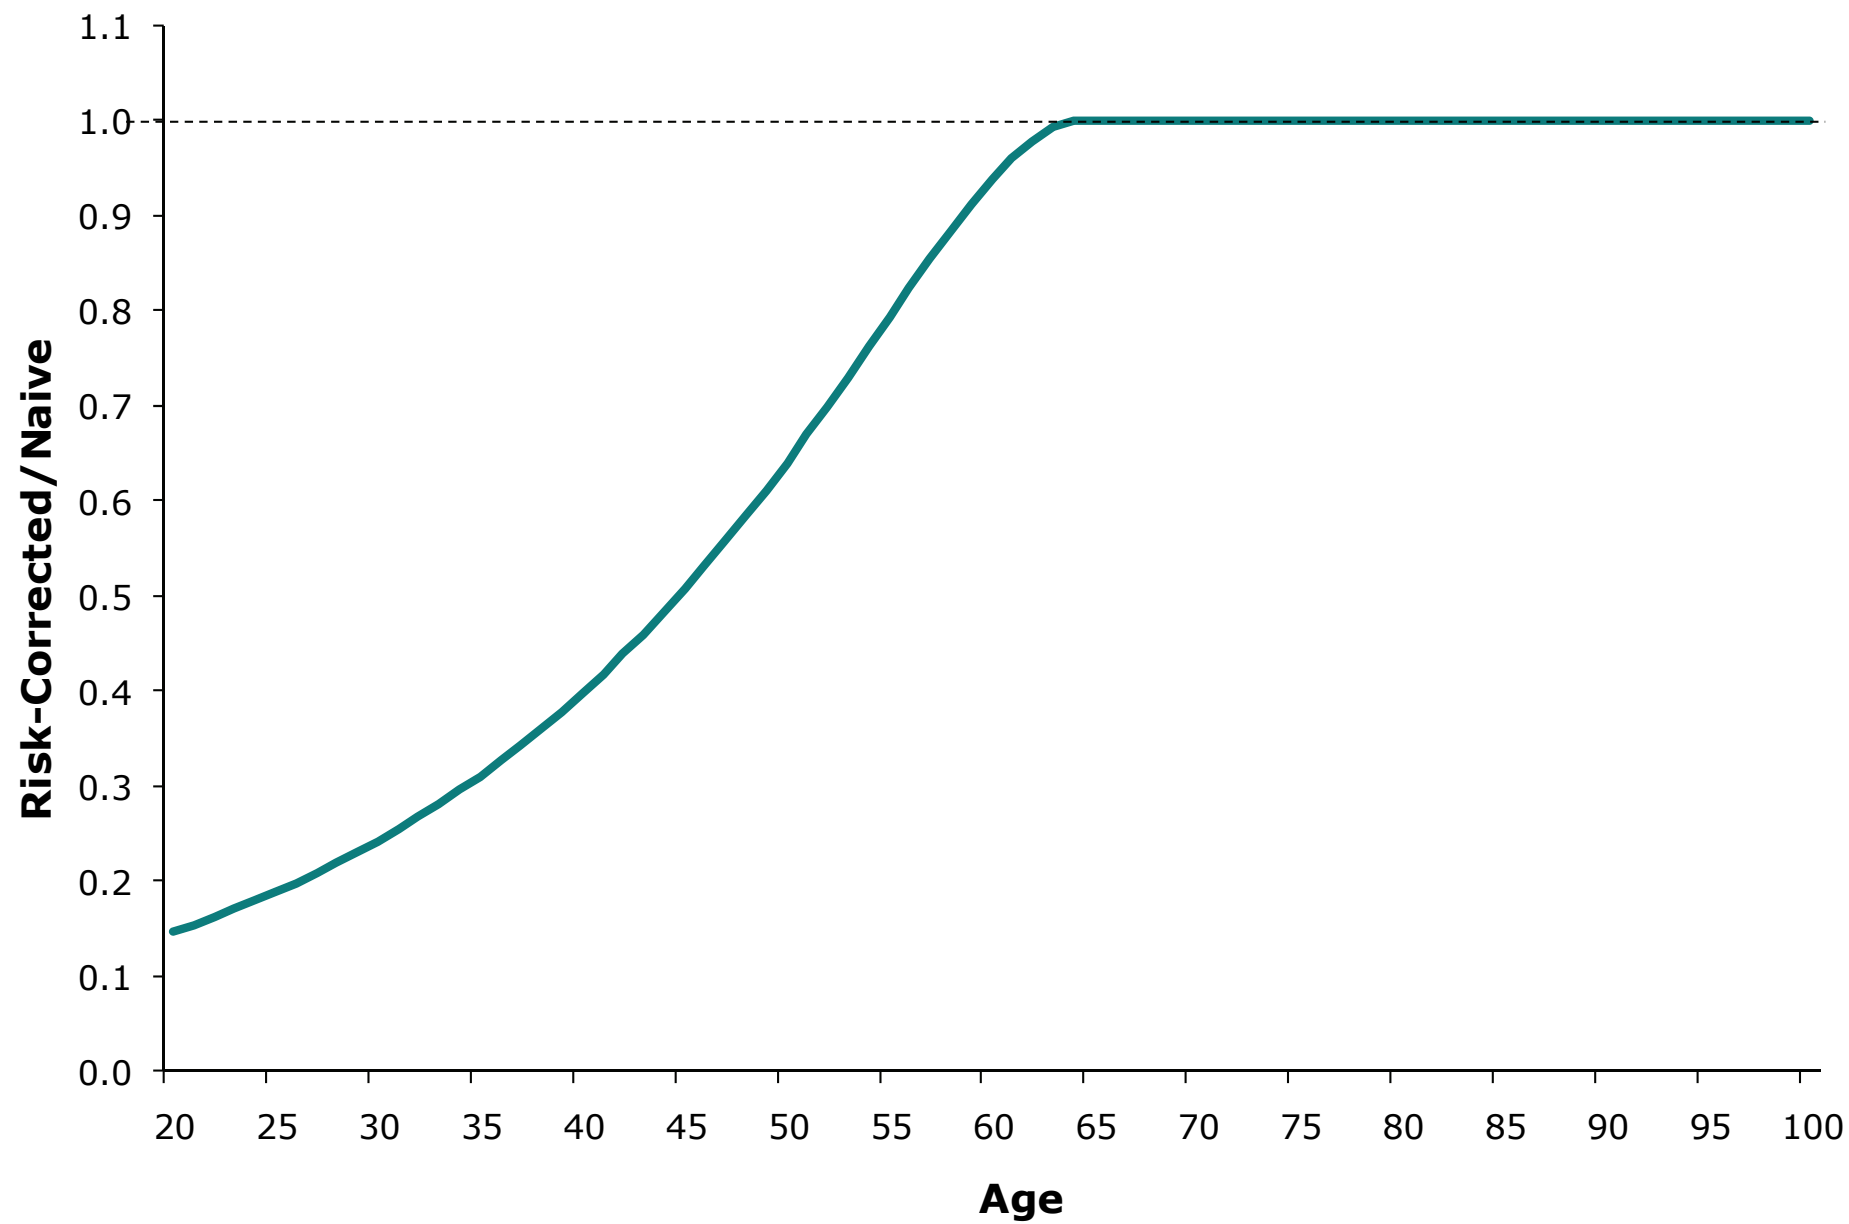**Figure 4: PAAW Price Ratios**

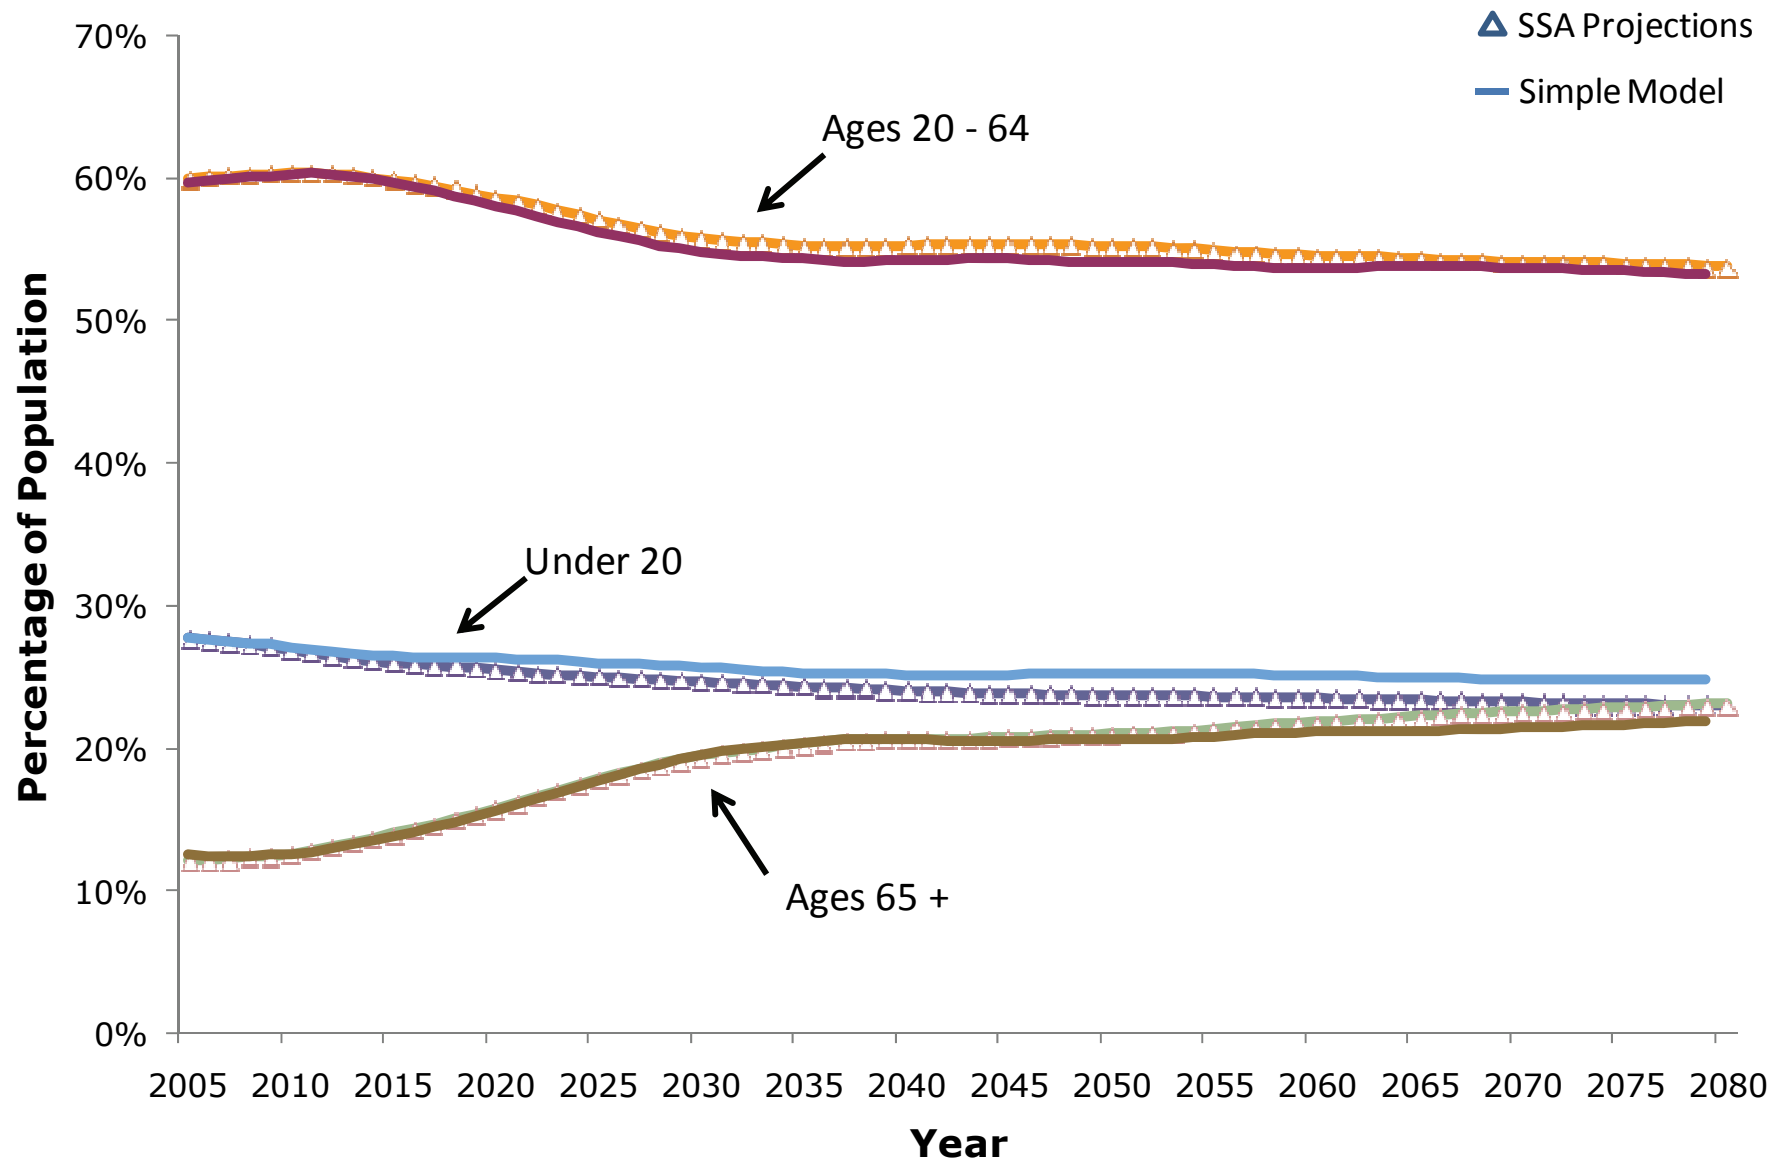**Figure 5: Projected Population Ratios, Model vs. Published**
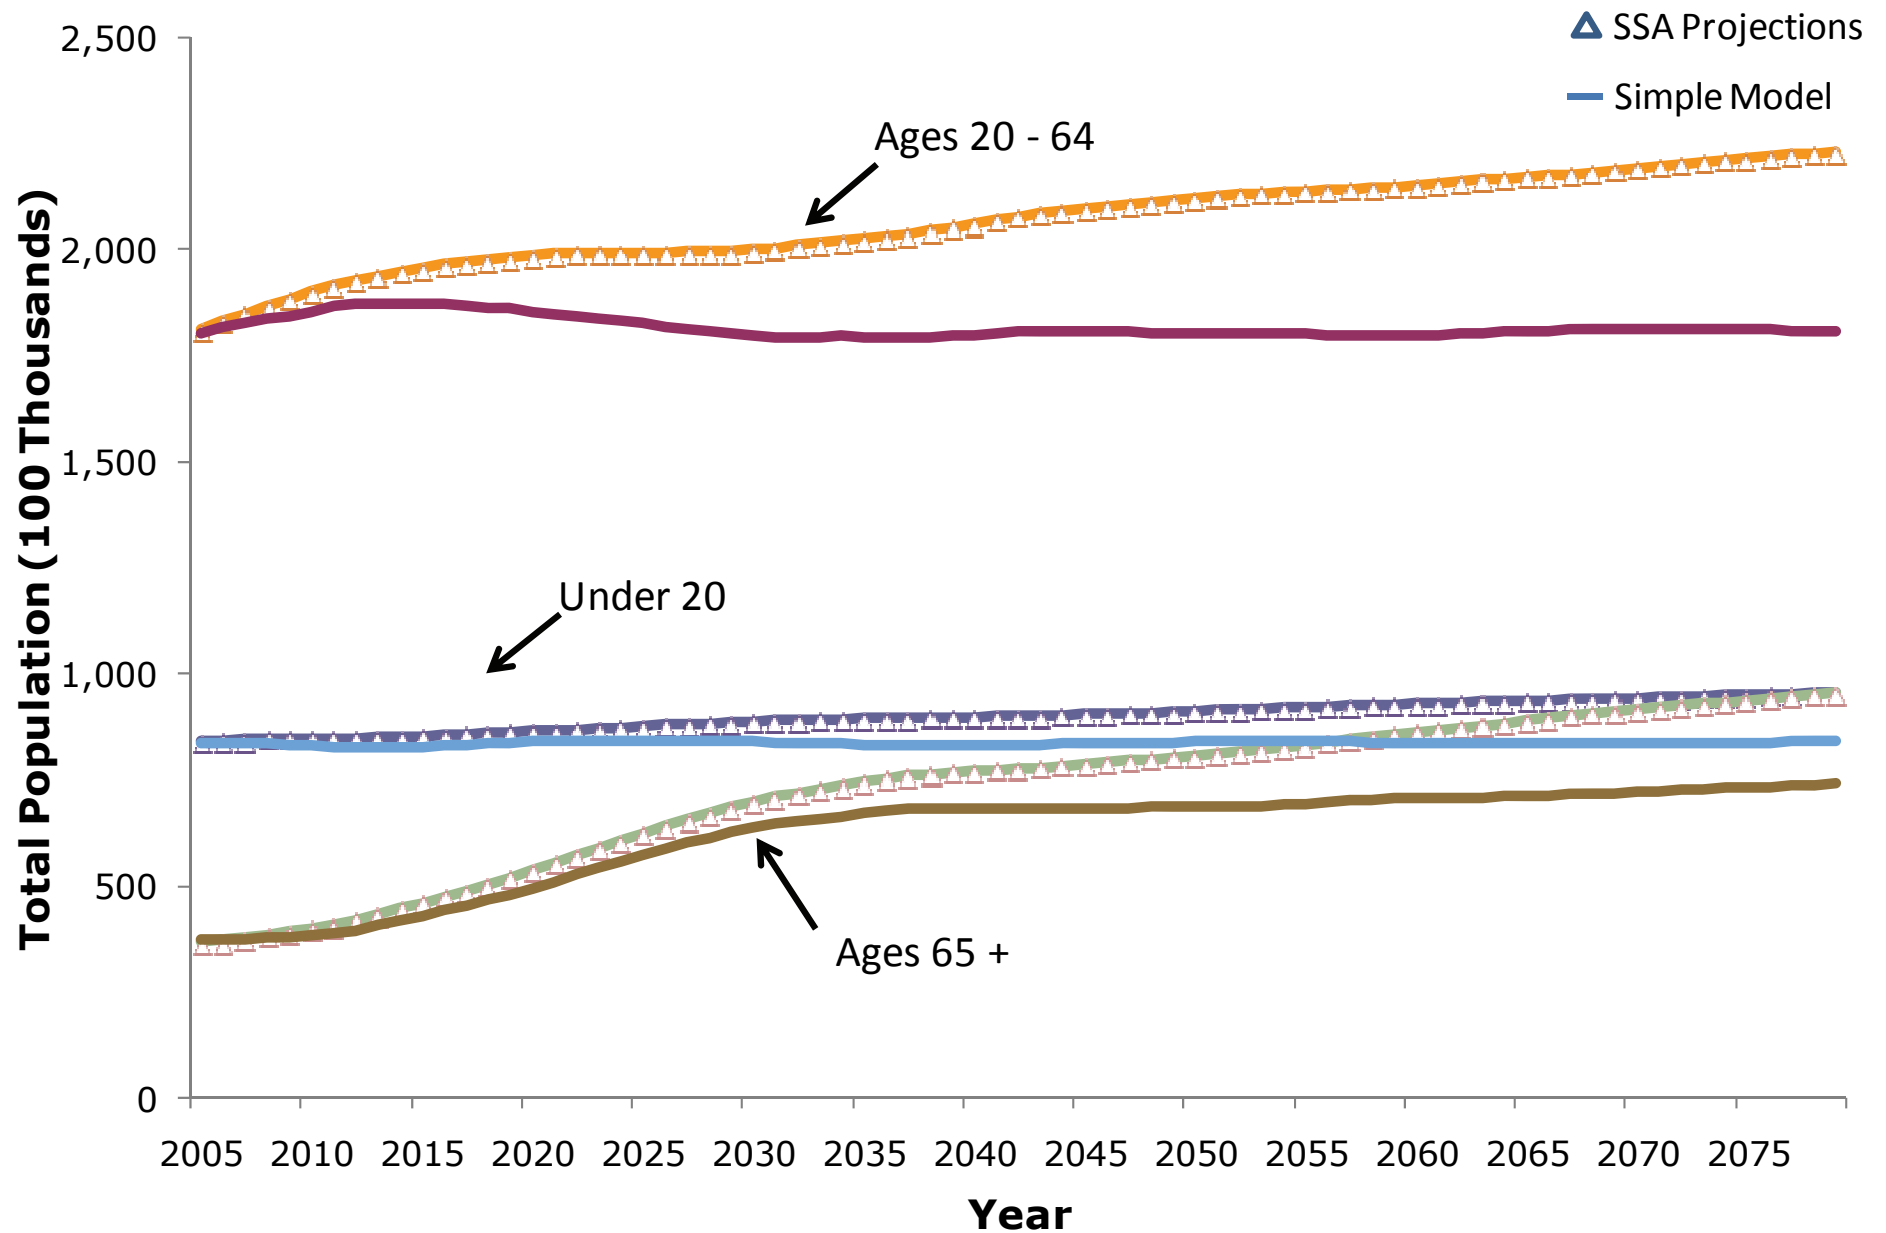
Figure 6: Projected Population, Model vs. Published
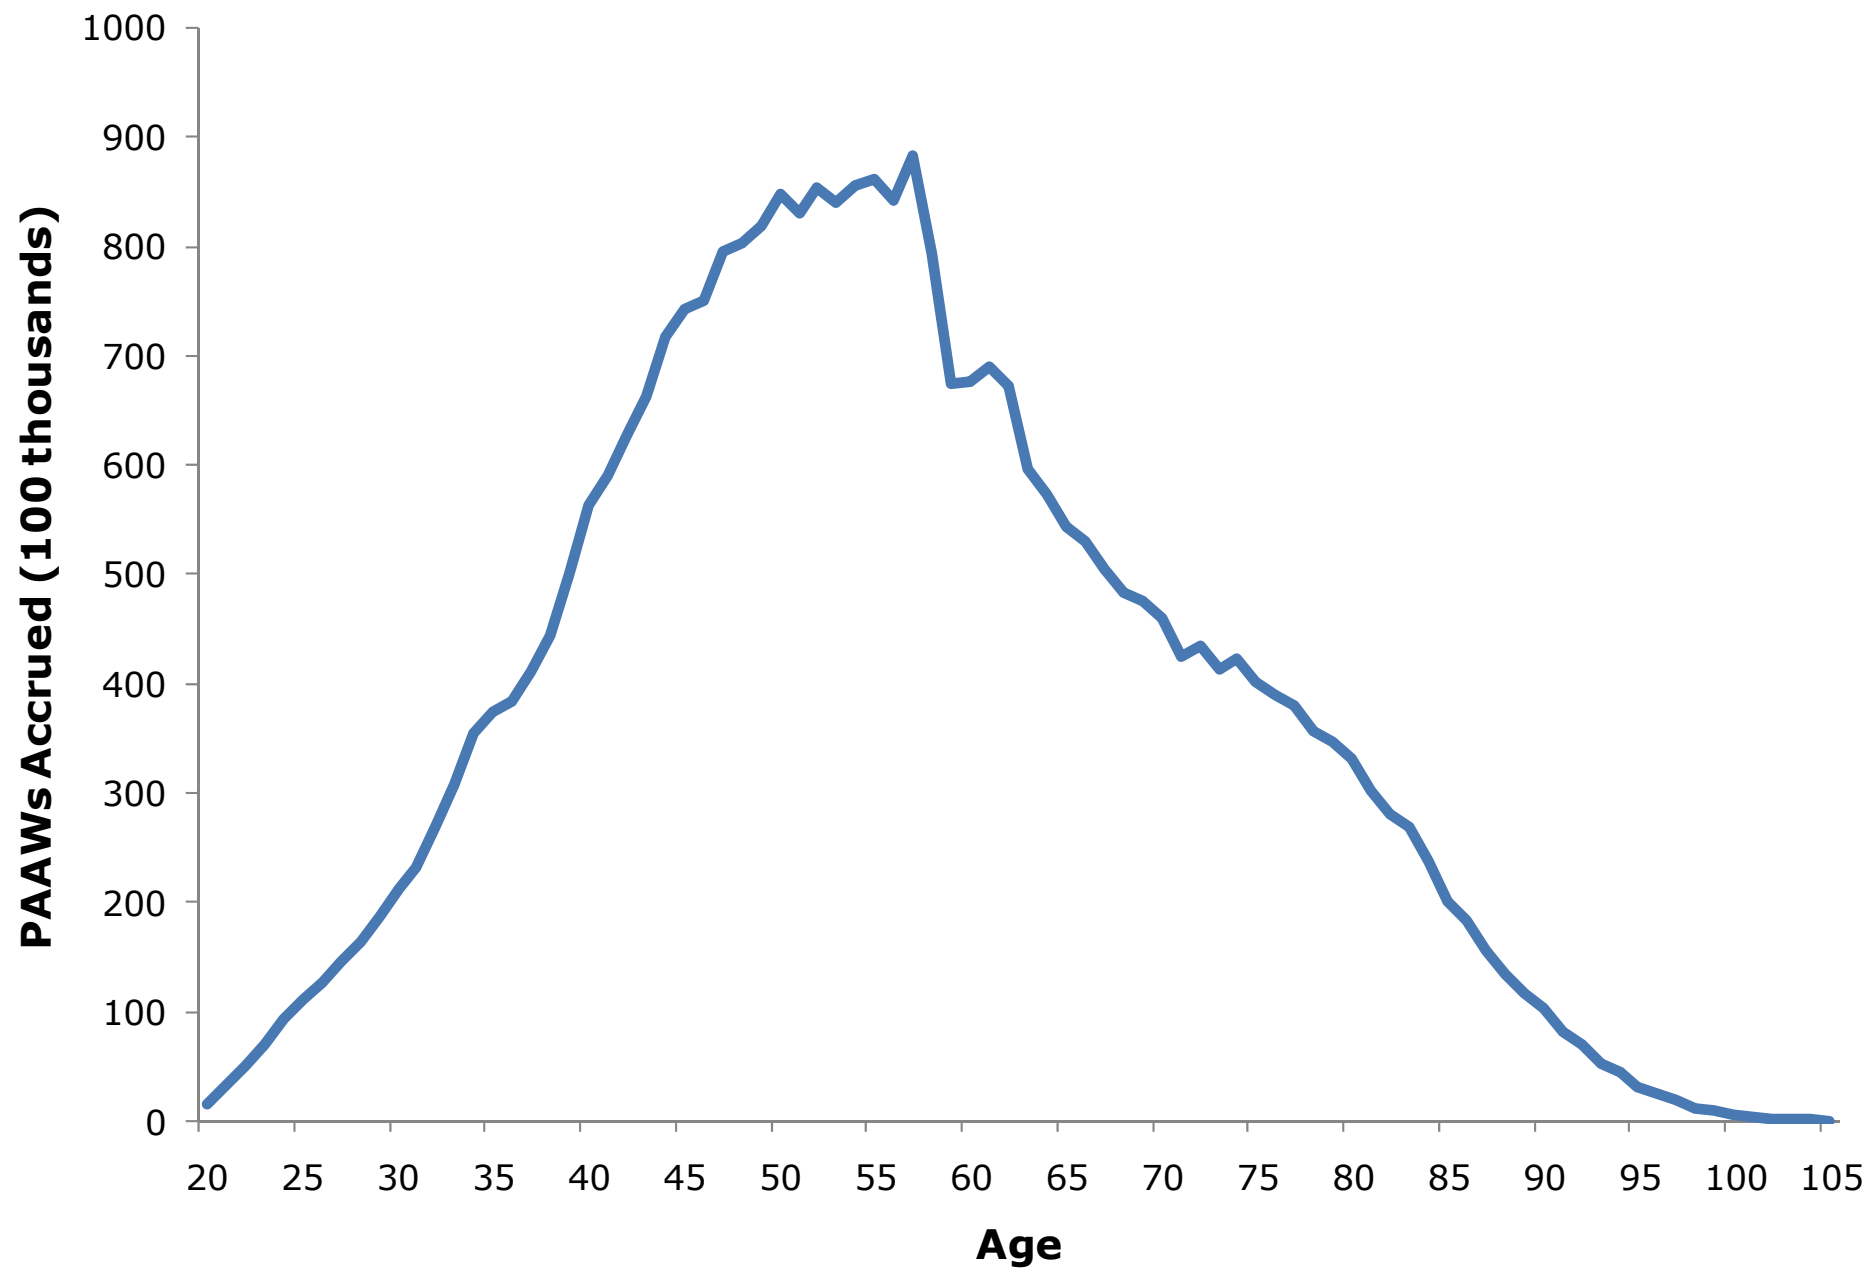
Figure 7: Cohort PAAWs Accruals

Figure 8: Cohort Value of PAAWs (OACT)

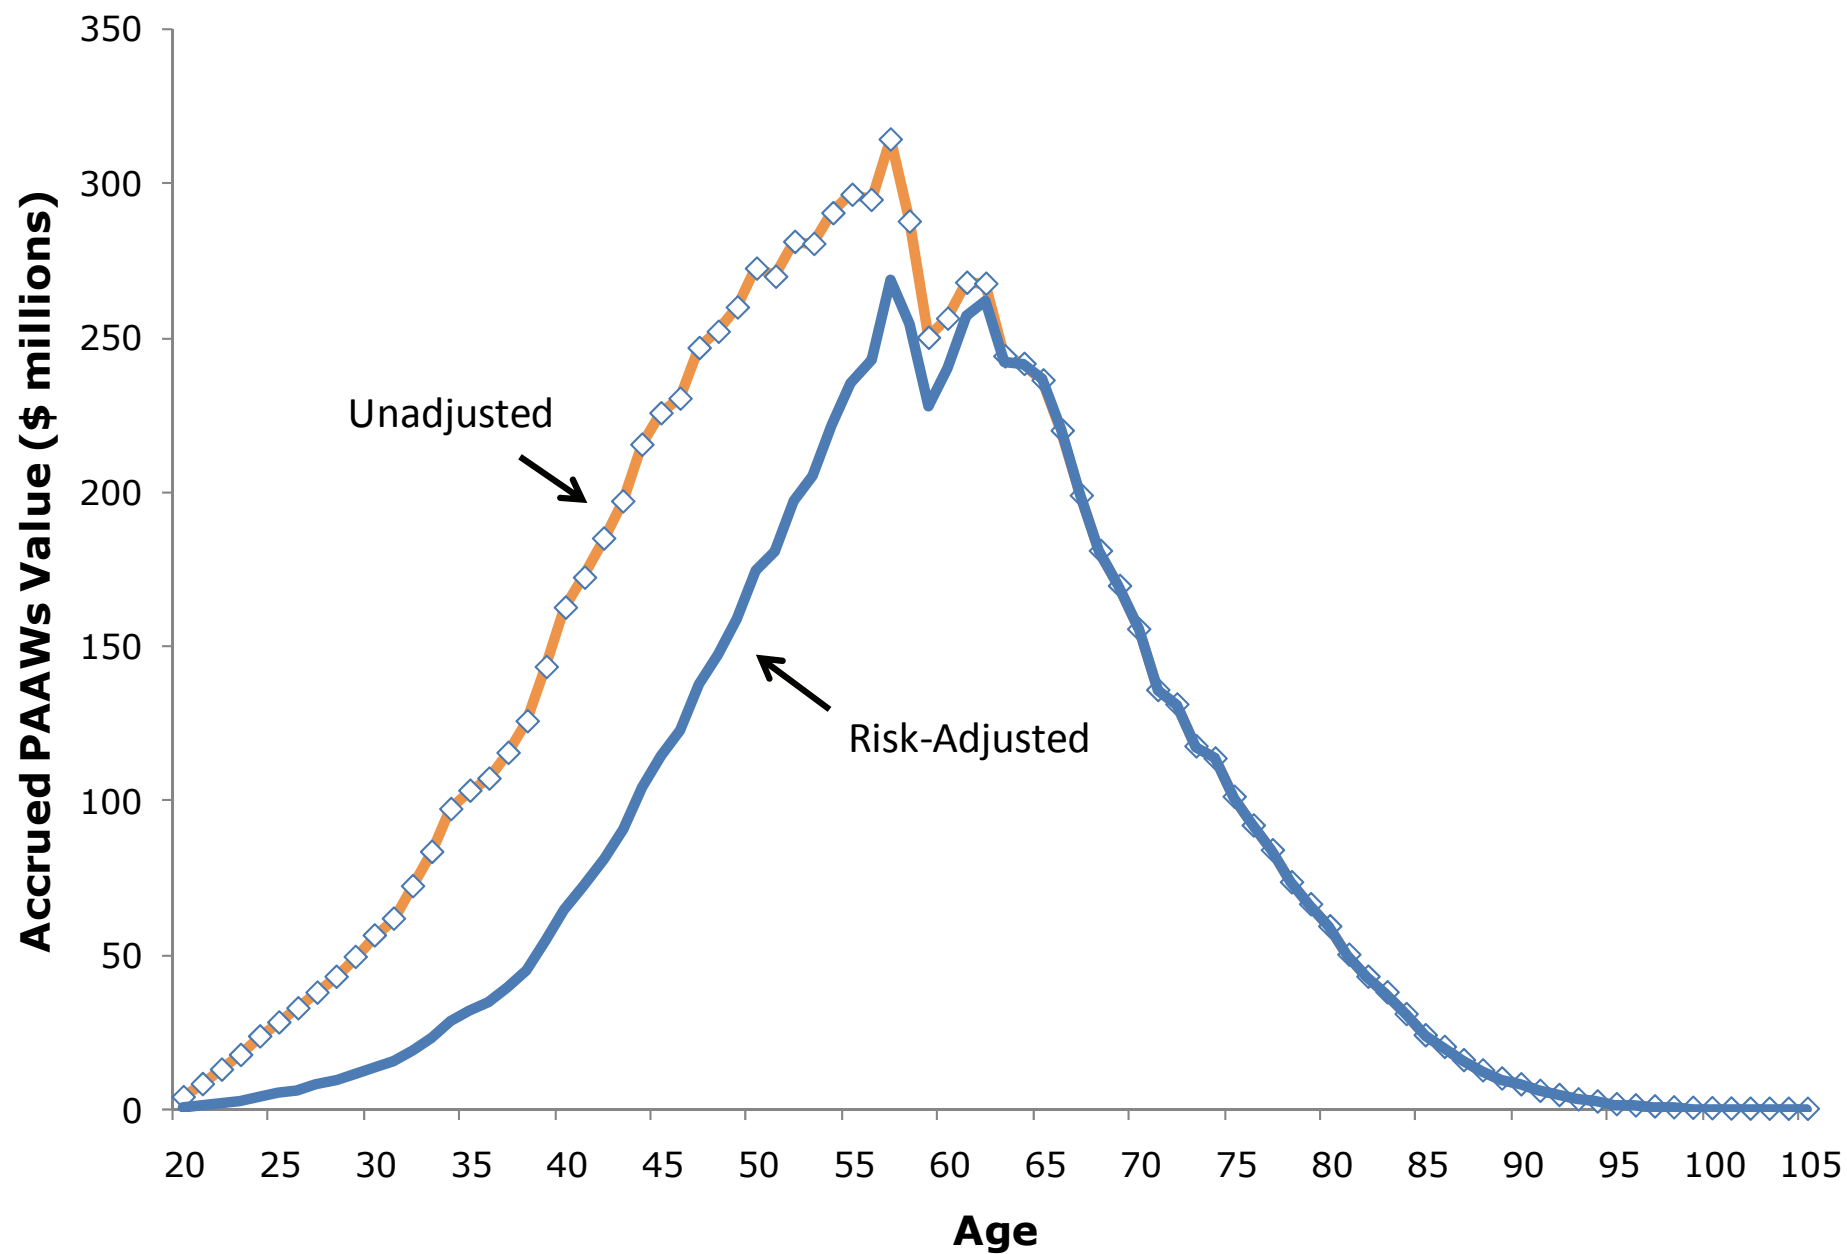**Figure 9: Projected Cash Flows, Model* vs. Official**

Legend:
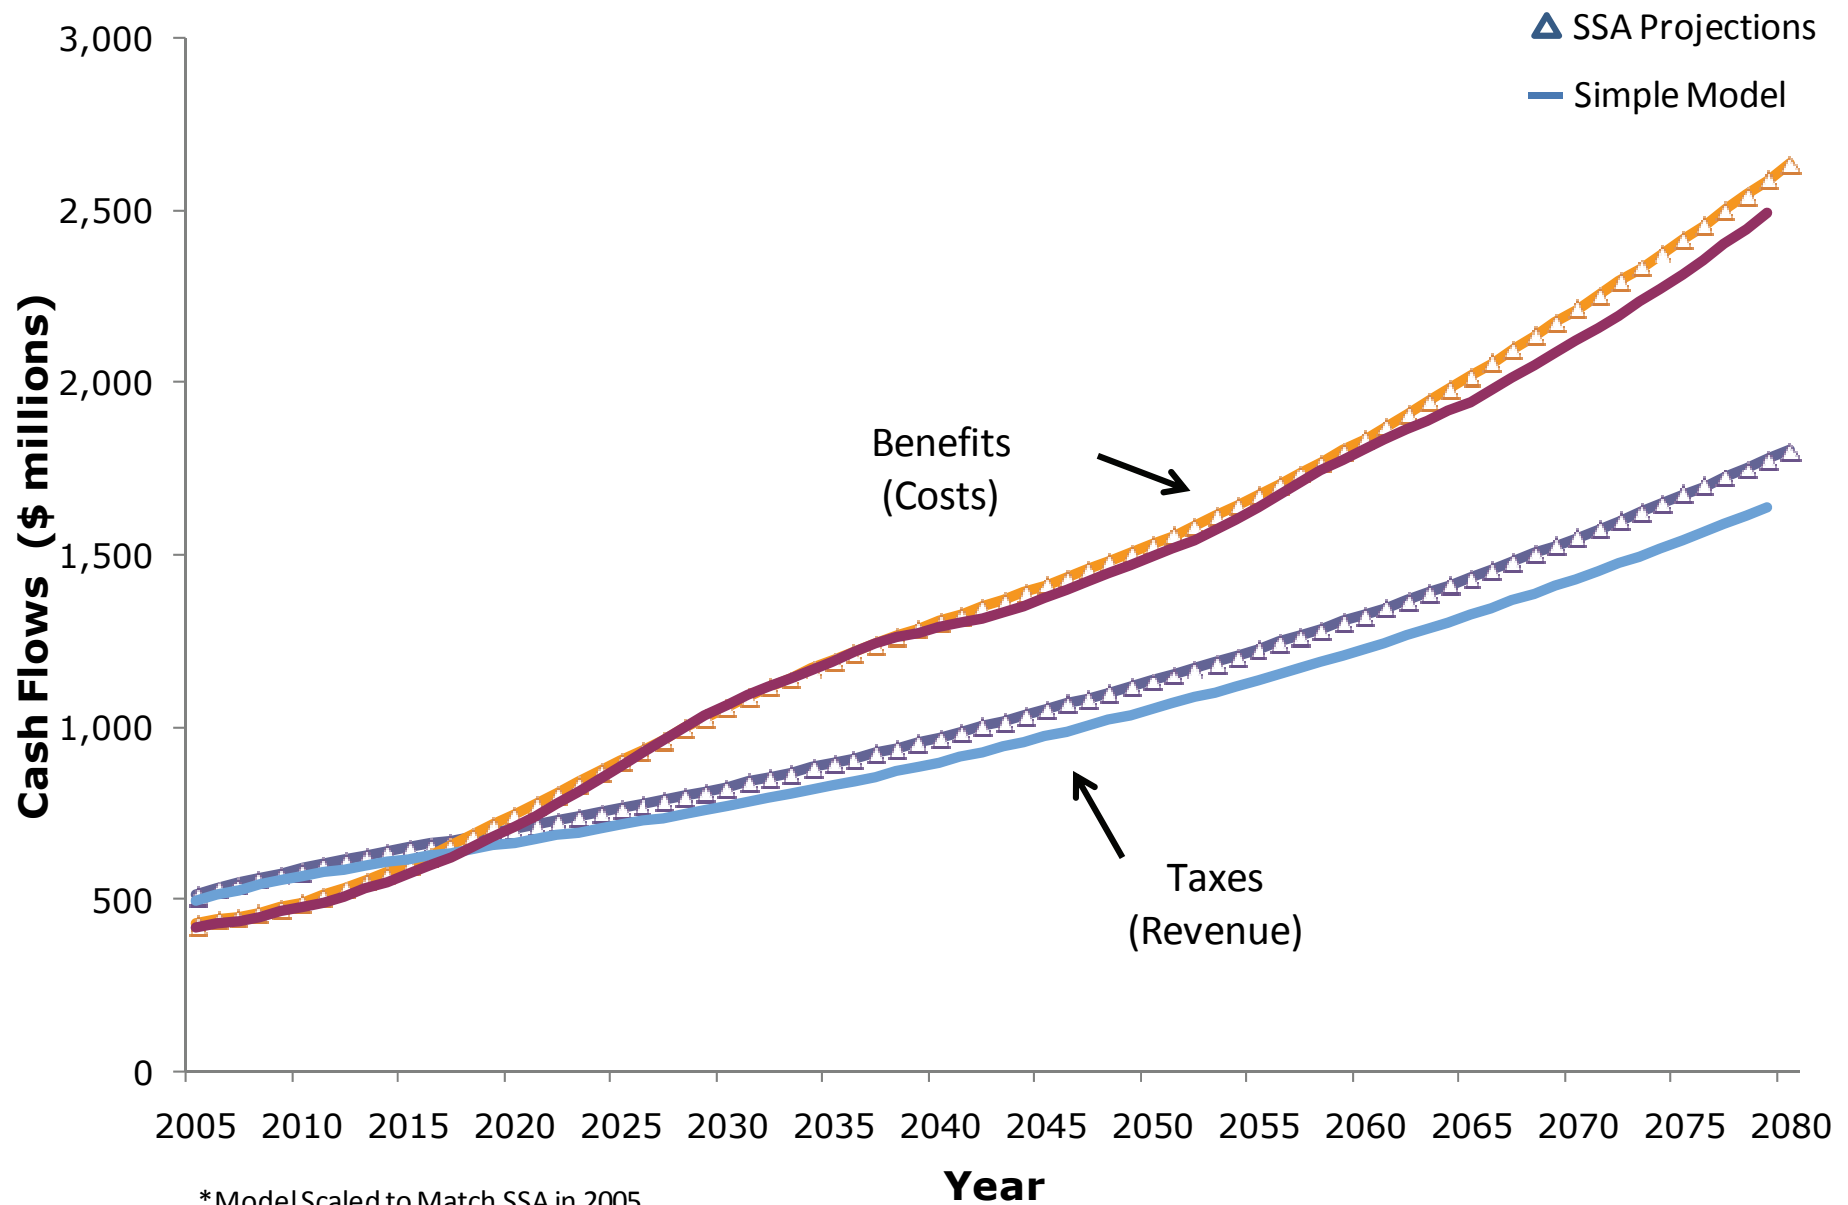- △ SSA Projections
- — Simple Model

Y-axis: **Cash Flows ($ millions)** — 0, 500, 1,000, 1,500, 2,000, 2,500, 3,000

X-axis: **Year** — 2005, 2010, 2015, 2020, 2025, 2030, 2035, 2040, 2045, 2050, 2055, 2060, 2065, 2070, 2075, 2080

Benefits (Costs)

Taxes (Revenue)

*Model Scaled to Match SSA in 2005

Figure 10: Projected PV Cash Flows, Risk-Adjusted Model vs. Unadjusted (Model)
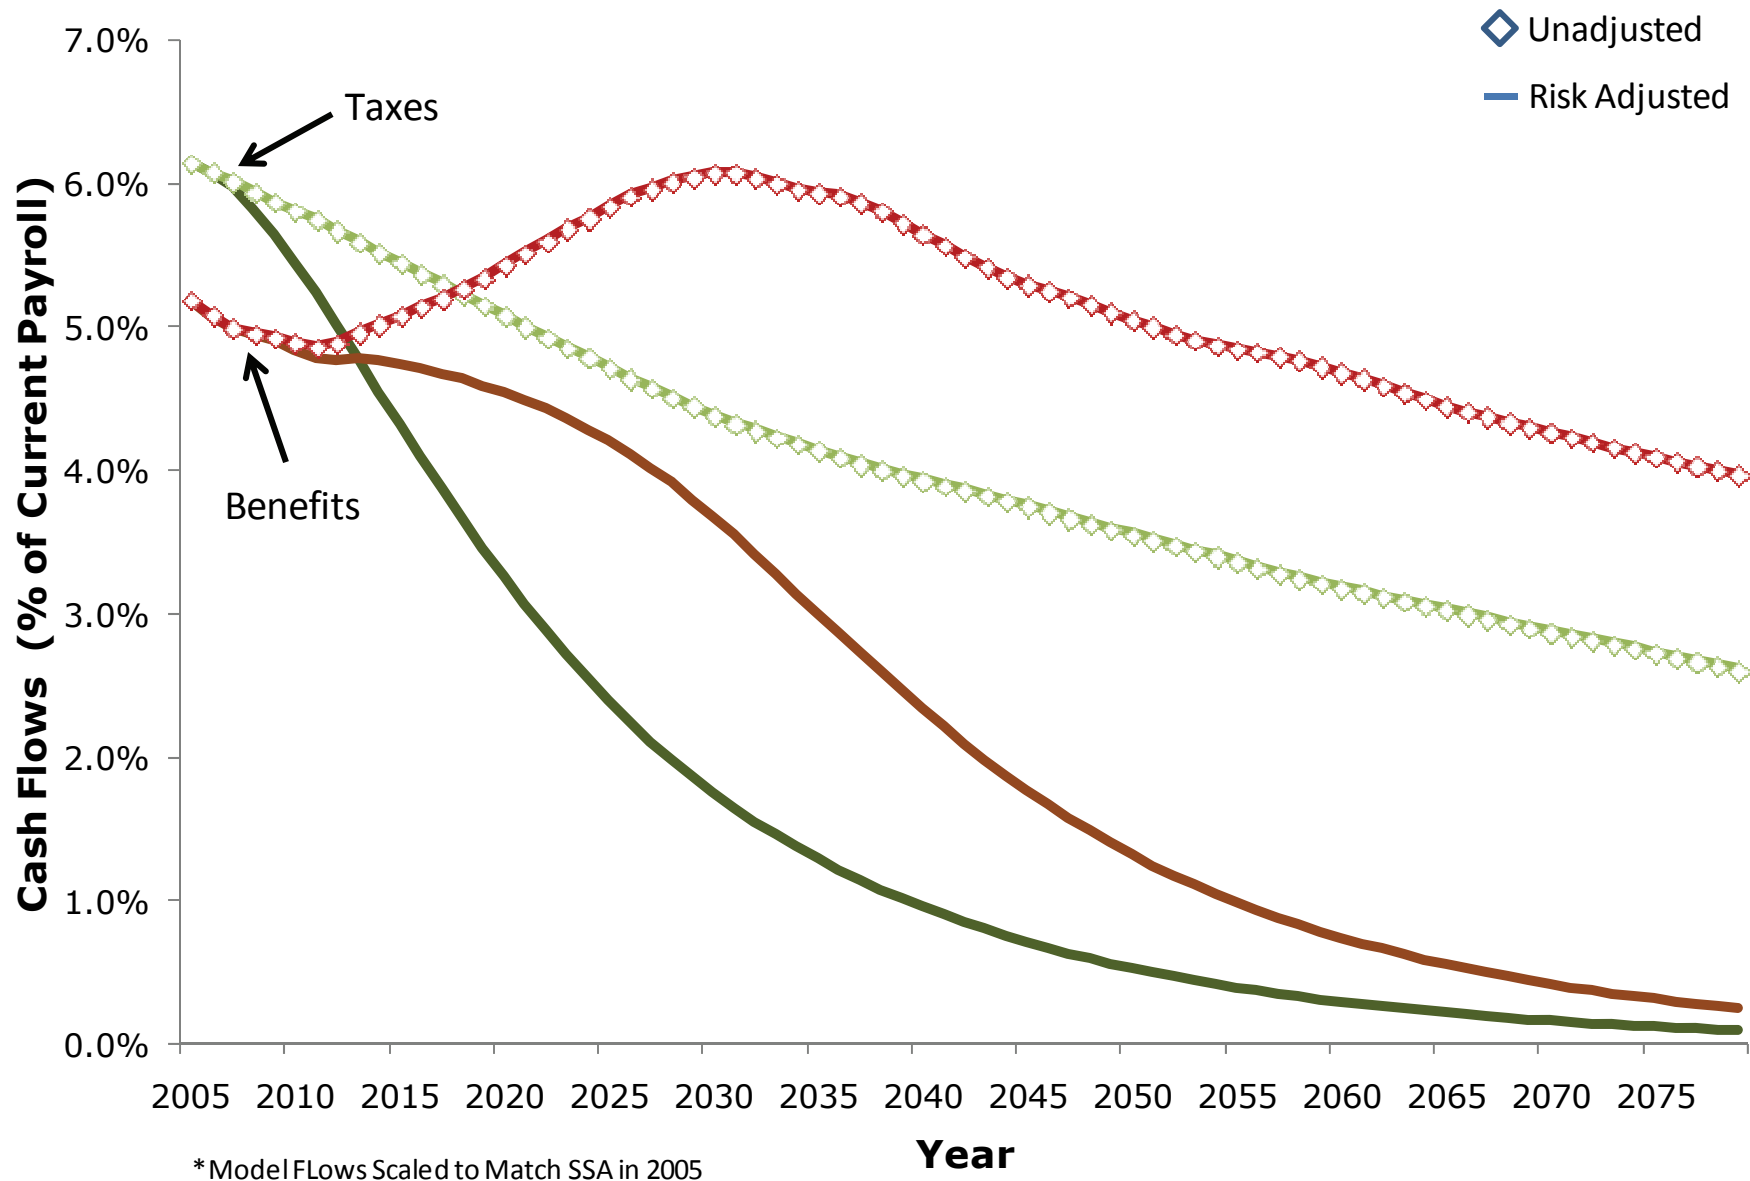
*Model FLows Scaled to Match SSA in 2005

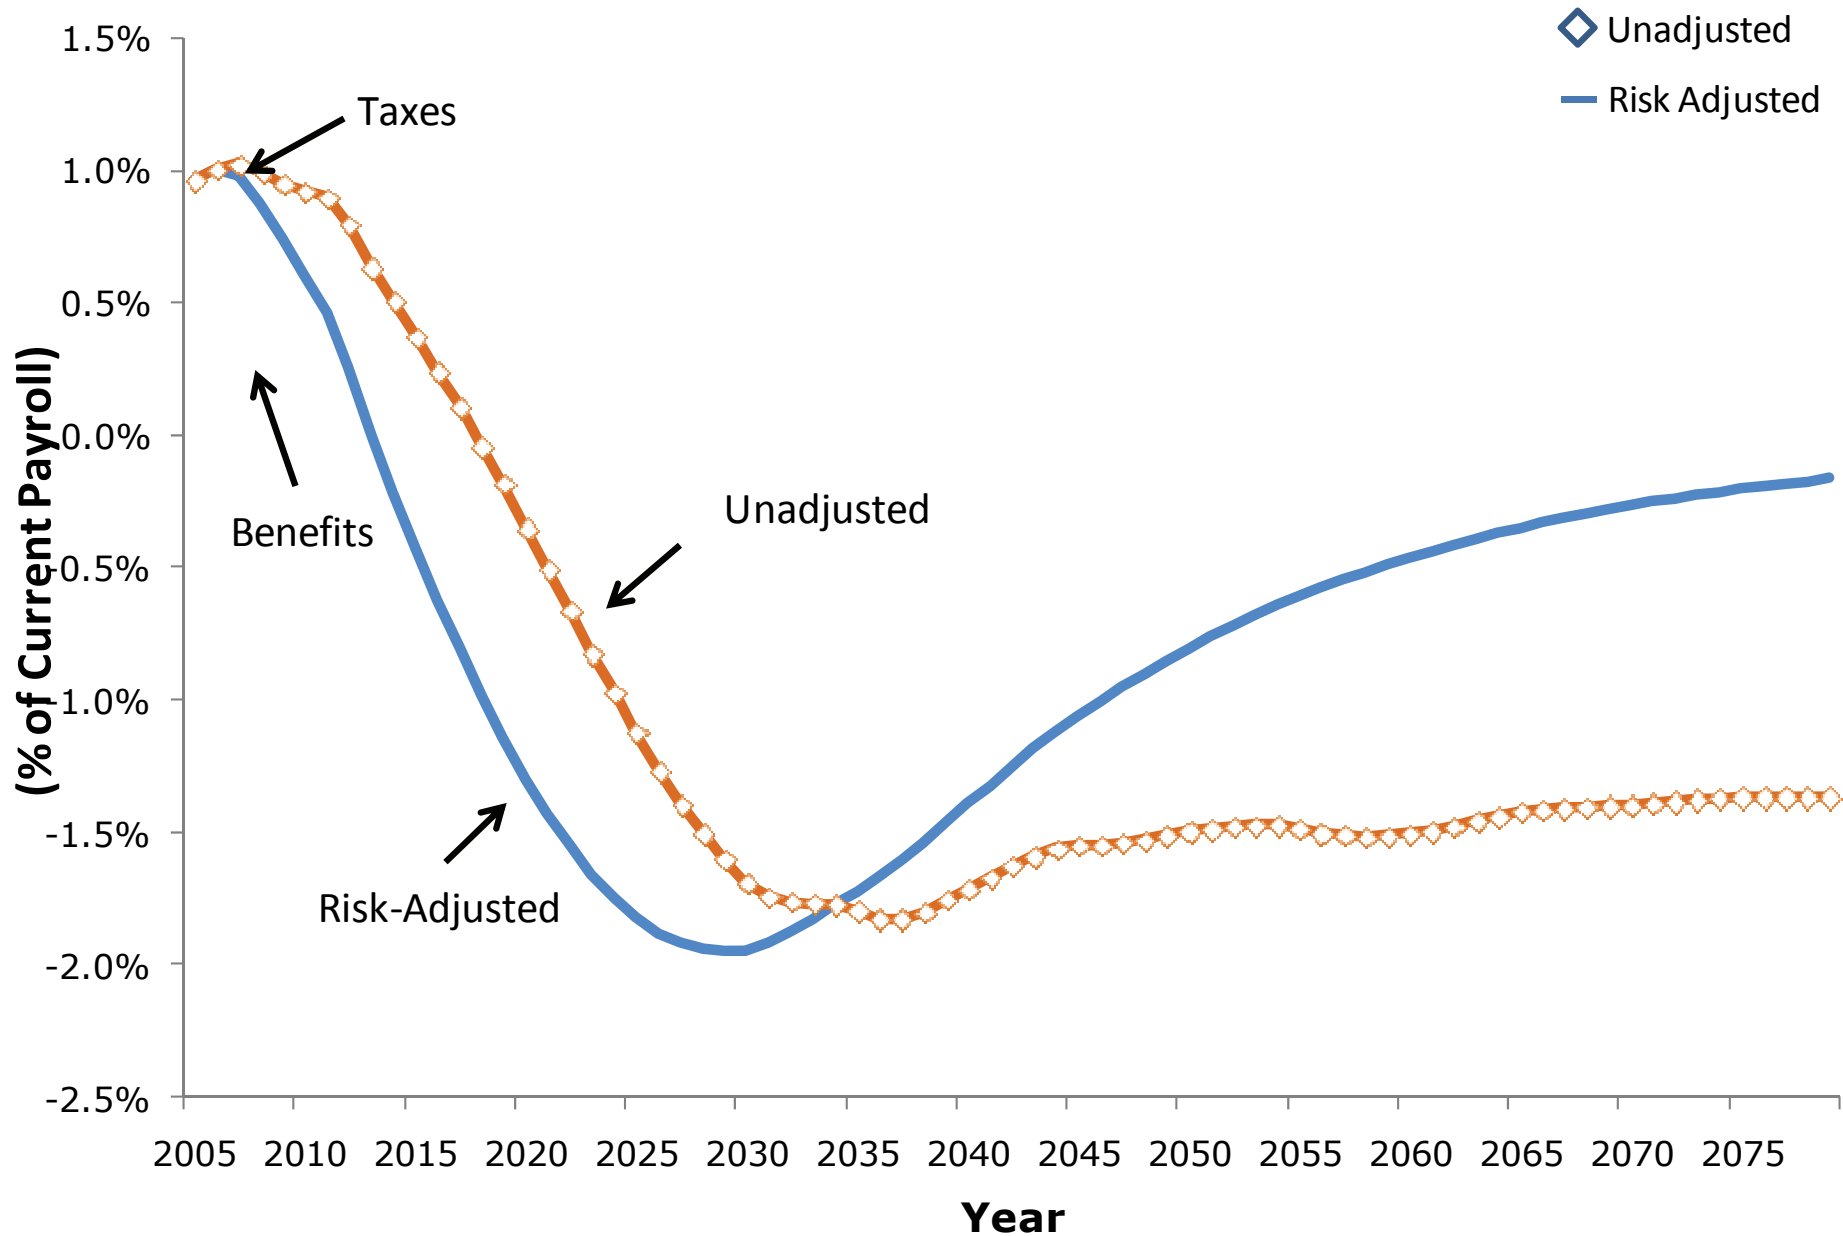Figure 11: Net Projected PV Cash Flows

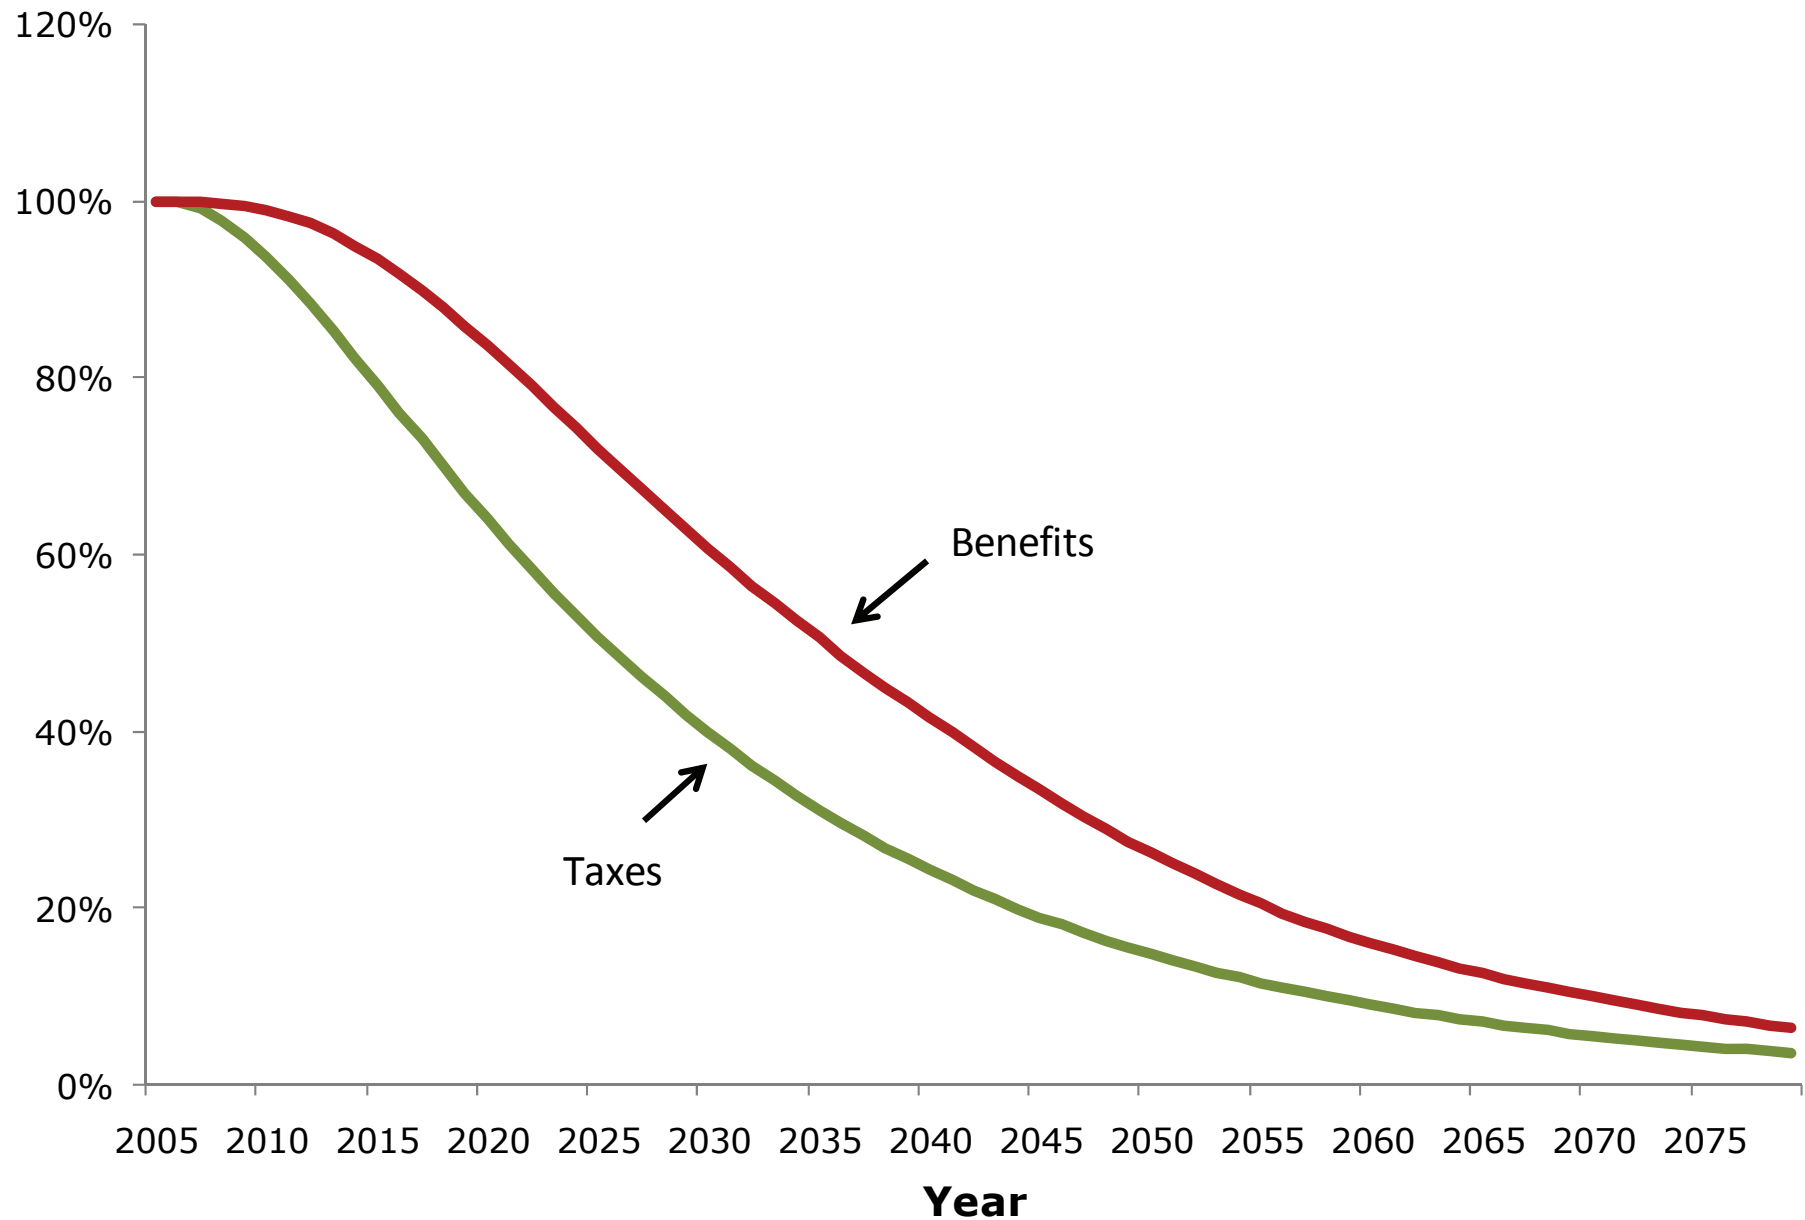**Figure 12: Ratio of Flows - Risk-Adjusted/Unadjusted**

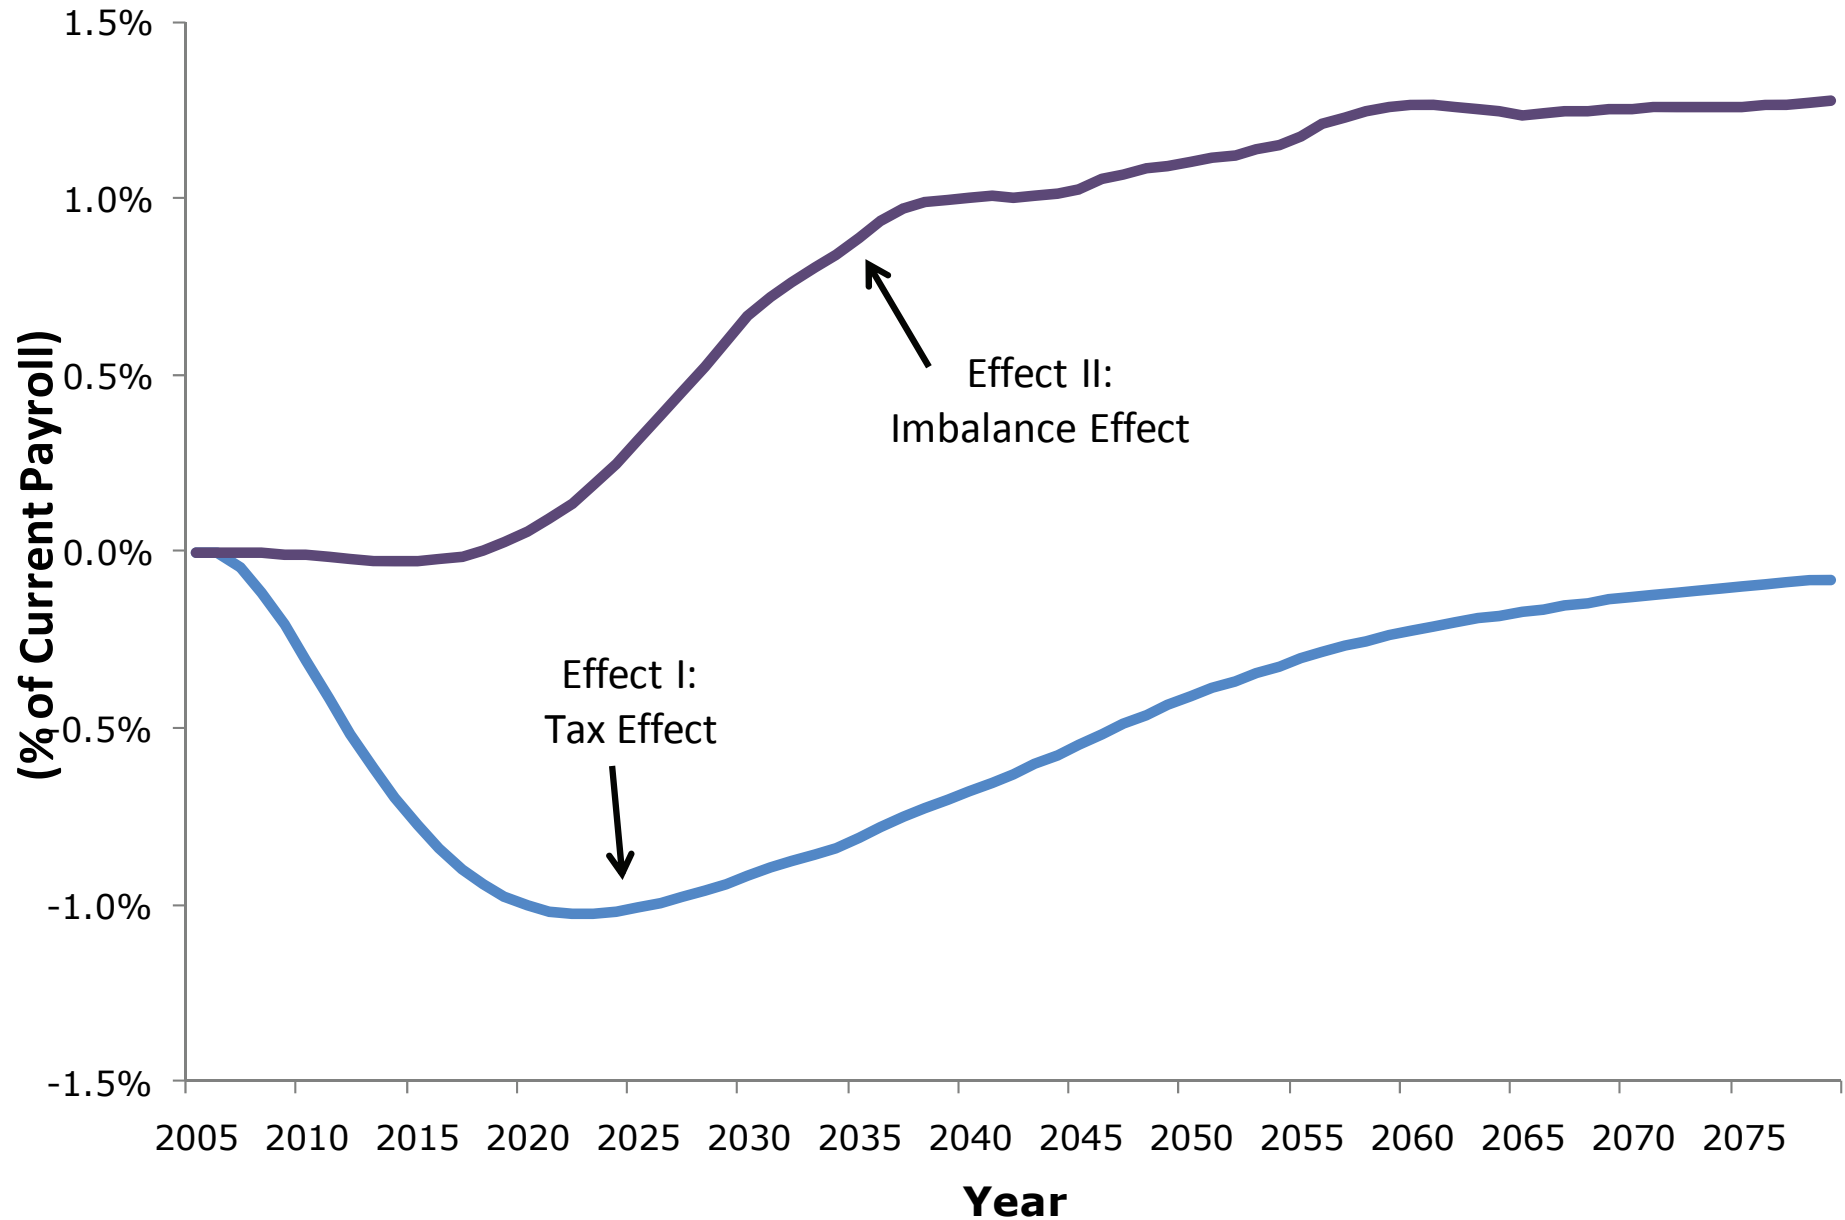# Figure 13: Two Different Effects of Risk Correction

Table 1: Components of Social Security Actuarial Balance (model)

(Percent of initial period payroll, excluding trust fund assets)

	Tax Revenue	Benefit Outlays	Total Surplus
Risk Corrected			
Open Group	2.1	3.1	-0.9
Closed Group	1.5	2.9	-1.4
Max Trans	-	2.2	-2.2
Naïve			
Open group	5.1	6.4	-1.3
Closed Group	2.1	5.0	-2.9
Max Trans	-	3.0	-3.0
Risk Corrected/Naïve			
Open group	0.42	0.48	0.72
Closed Group	0.71	0.59	0.50
Max Trans	-	0.74	0.74