

The Market Value of Social Security

by

John Geanakoplos

Yale University

and

Stephen P. Zeldes

Graduate School of Business, Columbia University; and NBER

Latest revision: July 11, 2011

We thank Ryan Chahrour and Ben Marx for research assistance, Deborah Lucas for helpful comments and suggestions, and staff at the Social Security Administration's Office of the Chief Actuary for help with data on projected contributions and benefits.

This research was supported by the U.S. Social Security Administration through a grant to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the author(s) and do not represent the views of SSA, any agency of the Federal Government, or the NBER.

The Market Value of Social Security

What discount rate should the government use to measure the financial status of various government programs and the impact of potential policy changes? How, if at all, should the discount rate take into account the risk of the underlying cash flows? We examine these questions in the context of one of the largest components of the U.S. budget: Social Security. Official measures of the U.S. Social Security system's present value funding shortfall are computed using the risk-free interest rate. Yet future Social Security expenses and revenues are, by law, linked directly to future realizations of the aggregate wage. We build a model in which consumption and wages are cointegrated and recompute common measures of Social Security's funding status using the model-implied risk adjustment. We find that the risk-adjusted shortfall in Social Security is between 18 and 68% less (depending on the measure) than indicated by the official statistics.

We also examine three policy changes that have been proposed to bring the system back into balance: increasing payroll taxes, decreasing benefits, and indexing benefits to prices rather than wages. For the first, we find that even though adjusting for risk reduces the present value shortfall, it actually increases the tax rate increase necessary to bring the system back into balance. For the second proposal, risk adjustment has ambiguous effects on the magnitude of benefit cuts. For the third proposal we find, contrary to the results of models that ignore risk, that linking benefits to prices instead of wages would *increase* the cost of Social Security benefits and thus *increase* the shortfall.

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 M. 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

What discount rate should the government use to measure the financial status of various government programs and the impact of potential policy changes? In practice, the government often uses a risk-free rate of interest to discount expected future expenditures and revenues. But if the relevant cash flows are risky, then this approach yields present values that are very different from those assigned by private actors in the economy. The results can be inaccurate and misleading.

How should the discount rate be adjusted to incorporate risk? When government programs consist of explicit or implicit holdings of securities that are actively traded in financial markets, the appropriate discount rate (or term structure of rates) equals the expected return(s) on these traded securities, and the corresponding present values are estimates of the “market value” of the program’s cash flows.¹ When the cash flows are not directly linked to market-traded securities, however, the best way to evaluate the market value of programs and policies is less obvious.

We examine these questions in the context of one of the largest pieces of the budget: the U.S. Social Security program. The cash flows in Social Security cannot be tied to any currently traded financial security, but nevertheless contain substantial market risk. We argue that risk adjustment is appropriate and we show how to compute the appropriate risk-adjusted discount rate and the corresponding market value estimates of the cash flows. We argue that the best way to incorporate a correction for risk is to estimate what the market value *would be* if claims on these flows were traded in financial markets. This concept corresponds to that of “fair value” in the accounting literature.

While there are many sources of demographic and economic risk in future Social Security cash flows, we focus on the source that we think is most important: wage risk. Social Security system rules tie both benefits and tax revenue to future economy-wide wages. We argue that in the long run, per capita wages, per capita consumption, and the value of the stock market are likely to be tightly correlated, in which case financial markets would add a risk-premium to the discount rate (or set of discount rates) used to value wage-indexed streams, decreasing the present value of the cash flows relative to

¹ For example, see the analyses of mortgage guarantees and student loans in Lucas (2010).

the “actuarial approach” that uses the risk-free rate for discounting. This adjustment decreases the value of *both* future benefits and future tax revenues, relative to the official estimates, and therefore the effect on the deficit depends on the net effect of the two adjustments.

In order to quantify the importance of this required risk adjustment, we calibrate two related models: 1) a structural stochastic real business cycle (RBC) model with advance “news” about future productivity and 2) a simple reduced form model of the link between stock prices and wages.

A key step in ascertaining the appropriate discount rate and estimating the market value of future benefits and taxes is determining the market price and expected return of a wage-linked security that we call a wage bond. We define a j -year wage bond as a security that pays, in period $t+j$, the economy-wide average labor earnings in that period. We use our models to compute the market yields on and prices of these wage bonds. We then use these prices to estimate market value measures of Social Security’s financial status and the costs of alternative policy reforms. We design these models so that they are consistent with the observations that wage growth and stock returns are essentially uncorrelated in the short run, but highly correlated over long horizons. Our work builds on that of Jermann (1999), Lucas and Zeldes (2006), Benzoni, Collin-Dufresne, and Goldstein (2007), and Geanakoplos and Zeldes (2010).

Our first model is based on a representative agent economy with Cobb-Douglas production and within-period utility given by a non-separable, concave function of consumption and leisure. In equilibrium, per capita wages and per capita consumption are always proportional to each other. We allow for stochastic productivity growth that augments labor and capital equally, and we calibrate stochastic output growth to match the volatility of annual output growth in the United States over the period 1947-2008.

The distinctive feature of our simple general equilibrium (GE) model is that it incorporates a “look ahead” component: we assume that agents receive news each period about productivity growth in subsequent periods. We calibrate the degree of look-ahead in the model by introducing a parameter (ρ) that measures how much agents know about future productivity growth: we assume that knowledge about the annual growth rate j periods into the future drops off at the rate ρ^j .

We compute analytical expressions for equilibrium wages, asset prices (including

the price of stocks and wage-bonds), and asset returns. News about the future can bring large swings in asset prices, even though in the model uncertainty about productivity and wages in period t is entirely resolved by period $t-1$. The contemporaneous correlation between wage growth and stock returns is therefore zero, despite the fact that the two are perfectly linked in the long run.

Next, we compare our risk-corrections to those generated by a general stochastic process with cointegrated wages and stock prices, of the form considered and estimated by Benzoni, Collin-Dufresne, and Goldstein (2007). Assuming that the variation in wages that is uncorrelated with stock prices generates no additional risk-premium, a standard application of modern financial tools (risk-neutral pricing) allows us to price wage bonds, and thus Social Security cash flows, as derivatives on the stock market. Previously, other authors have used processes like this to perform asset pricing exercises, including Lucas and Zeldes (2006), who used a similar process to value private defined-benefit pension liabilities when benefits are tied to the final salary of the individual employees, and Geanakoplos and Zeldes (2010), who used this process to estimate the market value of accrued Social Security benefits. Although this model is quite different conceptually from the general equilibrium model, we show that the set of risk corrections implied in this alternative model are identical to those implied by an appropriately calibrated version of the look-ahead general equilibrium model.

We define the risk-adjustment, RA , (a.k.a the “mark-down ratio”) as the ratio of the market (risk-adjusted) price of wage bonds to the actuarial (non-risk adjusted) price. We derive a simple analytical expression for this ratio, valid in both models, that depends only on the time horizon j , the parameter measuring the strength of the relationship between wages and stocks, and the risk characteristics of the economy. A convenient property of both models is that this mark-down ratio does not depend on expected wage growth or the level of the risk-free rate.

From the $RA(j)$, we derive the constant discount rate for each wage bond that, if used, would correctly value that bond. This term structure of discount rates is the appropriate one to use to discount future wage-linked cash flows. We show that the $RA(j)$'s are less than one, and therefore that the discount rates are always greater than the risk-free rate, i.e. they should include a positive risk premium to account for market risk. In addition, we show that the risk premium is not constant, but instead increasing

with the horizon of the bond, i.e. the annualized risk-premium built into the discount rate is larger for long-horizon cash flows than for short-horizon cash flows. This is because wages are much less correlated with market fundamentals in the near term than at longer horizons.

We argue that several authors (e.g. Goetzmann, 2008, and Blocker, Kotlikoff, and Ross, 2009) have erred in applying standard factor models of asset pricing to the problem of valuing wage-linked benefits, because they assumed that the absence of short-run correlation between stock returns and wage growth implied little market risk in long-dated wage-indexed cash flows.² Because our models imply no short-run correlation but a strong long-run correlation between wage growth and stock returns, we show that applying their techniques to data generated by our model leads to inaccurate conclusions.

Once we have derived the prices and implied discount rates for wage bonds, we use these prices to estimate the market value of four measures of Social Security's financial status: 1) the *immediate transition cost* (ITC), equal to accrued benefits to date minus the current value of the trust fund³; 2) the *closed group transition cost*, which adds in future tax contributions of current workers as well as the future benefits accrued as a result of those contributions, 3) the *open group unfunded obligation*, which adds in future taxes and benefits of all future workers (over, e.g., 75-year and infinite horizons) and 4) a *future transition cost* (FTC) measure equal to the finite horizon (e.g. 75 years) open group unfunded obligation plus the projected value of the transition cost at the end of the finite horizon. To estimate the market value of these measures, we first convert the tax and benefit cash flow projections used in the SSA Trustee's Report into units of wage bonds.⁴ For taxes, this is relatively straightforward, because tax revenue in each year depends directly on average wages in that year. For benefits, we incorporate the fact that a person's benefits in any future year depend on the average wages in the

² See Geanakoplos and Zeldes (2010) for a more detailed critique of Blocker, Kotlikoff, and Ross (2009). Koehler and Kotlikoff (2009) responded to this critique by estimating a separate regression for cash flows at each horizon. We will argue that their revised approach does not adequately account for the risk in long-dated wage-linked securities.

³ The Social Security Office of the Chief Actuary refers to this measure as the "Maximum Transition Cost" measure. See Wade et. al (2009).

⁴ In doing so, we ignore any uncertainty in the number of workers and beneficiaries.

economy in the year that the person reaches 60.

What are the effects of risk adjustment on the various measures? Because the *immediate transition cost* includes future benefit payments only, the market value shortfall must be less than actuarial value, and we estimate it to be 18% lower (\$11.7 trillion versus \$14.2 trillion). The situation is more subtle when it comes to valuing the open and closed group measures that include both future benefits and future taxes, as both are worth less as a result of risk adjustment. It turns out that the net effect is almost always that the shortfalls are smaller. For example, we estimate that the closed group shortfall is 41% smaller than the actuarial value (\$8.1 trillion versus \$13.7 trillion), and the infinite-horizon open group shortfall is 68% smaller than the actuarial value (\$6.4 trillion versus \$19.9 trillion).

It is worth noting that these results are reversed, and measured shortfall actually increases, for another prominent measure, the 75-year open group measure. We argue, however, that this measure does not correspond to any relevant conceptual experiment, and is therefore not particularly useful. Instead, we propose a new “future transition cost” measure (described above) and show that the market value shortfall is 51% smaller than the actuarial value (\$6.4 trillion versus \$13.1 trillion).

We next examine three policy experiments. The first involves raising tax rates by a flat percentage amount sufficient to bring the system into balance. Even though the present value of the shortfall is less under the market value approach, we find that because the present value of taxes is also less, the tax increase necessary to balance the system turns out to be *greater* than under the traditional actuarial approach. This is true for every imbalance measure when we consider a permanent increase and for all but one measure when we consider measure-specific payroll tax increases.

The second policy experiment is cutting benefits: we examine the percentage reduction in benefits that would balance the system according to each measure. For short-term measures the market adjustment works in the same direction as for the total amount of the imbalance, with smaller deficits requiring smaller benefit cuts to balance. The measures with longer horizons require more severe risk adjustments to the value of benefits, and for these measures the necessary benefit reduction is greater under market valuation even though the total imbalance is smaller.

The third policy experiment consists of shifting from the current system in which benefits are indexed to aggregate wages to one in which benefits are indexed to prices. The standard reasoning supporting this policy is that, since real wages have risen on average historically and are expected to do so in the future, indexing to prices should reduce future benefit outlays. While this is true on an expected value basis, on a risk-adjusted basis we find the opposite: namely that linking benefits to a price index instead of to a wage index *raises* the costs of future benefits and *increases* the shortfall. The intuition is that there are some future states of the world in which real wages are substantially lower in the future than they are today (i.e. wages will have risen less than prices), and in those states resources are particularly scarce and valuable, making it particularly costly for the government to pay out the higher price-indexed benefits. In our model, this effect is strong enough to outweigh the fact that in most of the other states wages will have risen more than prices.

While our paper focuses on Social Security, it has much broader public finance implications. Specifically, our results suggest that all government programs with costs and/or revenues tied to the performance of the macro-economy should be evaluated on a risk-adjusted basis.

Our paper is structured as follows. In Section II, we describe our two models and derive analytical solutions for the ratio of market to actuarial prices of wage bonds and the corresponding risk-premia that should be incorporated into the term structure of discount rates. We also summarize empirical evidence on the long-run relationships between wages and other macro and financial variables, and describe how we calibrate the model. In Section III, we describe the Social Security rules on taxes and benefits, show how to translate Social Security cash flows into units of wage bonds, and describe how we compute the market value estimates of the measures of Social Security unfunded obligations. Section IV presents our estimates of the market value of Social Security. In Section V, we examine our policy experiments. Section VI concludes.

II. Pricing wage bonds

Define a wage bond, $WB(j)$, as a security that makes a single payment j -periods from now equal to the average economy-wide wage in that period. There thus exists at time t a (potential) wage bond for each future period $t+1, t+2, \dots$, each of which has its

own risk-profile, sequence of returns, and price. In this section, we present two models for pricing wage bonds. The first is a simple general equilibrium macro model, and the second is a reduced-form asset pricing model. In both cases, we model the relationship between real variables and assume that inflation does not affect the relationship between real wages and real dividends.

A. A simple general equilibrium model

In this sub-section, we present a simple general equilibrium macroeconomic model in which productivity shocks represent the primary source of long-run macroeconomic fluctuations (and augment capital and labor productivity equally), and in which there is advance information (“news”) about future productivity changes that is not reflected in current productivity.

In the standard neo-classical growth model (and similarly in basic New Keynesian models), asset prices fully reflect agent’s expectations of future dividends (and thus future productivity), while wages are determined by the current level of productivity. The introduction of a significant “news” component generates fluctuations in asset prices that are not related to the current realization of productivity, and thus are also unrelated to current wages. When future productivity depends almost entirely on news shocks, wages and stock prices become contemporaneously orthogonal, without changing the long-run cointegration of the two that exists in the standard setup. In other words, this model generates the results that i) there is a positive long-run correlation between wages and stock prices, and ii) there is little or no short-run correlation between wage growth and stock returns.

Assume a neoclassical production function, $Y_t = \bar{A}_t K^\alpha H_t^{1-\alpha}$, where K represents a time-invariant capital stock, H represents hours per capita, and \bar{A} equals total factor productivity. Because capital is assume fixed, we can redefine $A_t \equiv \bar{A}_t K^\alpha$ and write the production function as $Y_t = A_t H_t^{1-\alpha}$. Assume firms own capital, produce, and sell output in

a perfectly competitive market. Let consumers have utility function $\frac{(C_t(1-H_t)^\theta)^{1-\gamma}}{1-\gamma}$. A

convenient feature of this utility function is that the disutility of labor is proportional to the

utility of current period consumption, which generates the result (shown below) that labor supply is constant. Note that because there is no investment, $Y_t = C_t$. Since markets are competitive, we can write the decentralized model as a social planner's problem. Necessary and sufficient equilibrium conditions can then be derived from the following Lagrangian

$$\max_{C_t, H_t} E_t \sum_{j=0}^{\infty} \beta^j \left[\frac{(C_{t+j}(1-H_{t+j})^\theta)^{1-\gamma}}{1-\gamma} + \lambda_{t+j} [A_{t+j}H_{t+j}^{1-\alpha} - C_{t+j}] \right]$$

Differentiating with respect to C_t and H_t yields (respectively) the following first order conditions:

$$\lambda_t = (C_t(1-H_t)^\theta)^{-\gamma} (1-H_t)^\theta \quad (1)$$

$$\lambda_t A_t (1-\alpha) H_t^{-\alpha} = (C_t(1-H_t)^\theta)^{-\gamma} C_t \theta (1-H_t)^{\theta-1} \quad (2)$$

Combining these equations with the constraint that $C_t = A_t H_t^{1-\alpha}$ gives the result that

$H_t = \frac{(1-\alpha)}{\theta+1-\alpha}$, $\forall t$. Call this this value H . Using equation (1) and the the fact that, under

a competitive decentralization, the wage equals the marginal product of labor, we have

$$W_t = A_t (1-\alpha) H^{-\alpha} \quad (3)$$

$$\lambda_t = A_t^{-\gamma} H^{-\gamma(1-\alpha)} (1-H)^{\theta(1-\gamma)} \quad (4)$$

All of the endogenous variables are thus defined in terms of the generic exogenous process A_t .

We now proceed to price a variety of assets, which need not, in general, trade in equilibrium. First, define the value of the firm as the present value of future rents accruing to capital owners (firm owners) discounted using the stochastic discount factor

$$\beta^j \frac{\lambda_{t+j}}{\lambda_t} :$$

$$V_t^{firm} \equiv E_t \sum_{j=1}^{\infty} \beta^j \frac{\lambda_{t+j}}{\lambda_t} (Y_{t+j} - W_{t+j} H_{t+j}) \quad (5)$$

$$= \frac{\alpha H^{1-\alpha}}{A_t^{-\gamma}} E_t \sum_{j=1}^{\infty} \beta^j A_{t+j}^{1-\gamma} \quad (6)$$

Second, consider a zero-coupon bond that pays one unit of the consumption good in period $t+j$. This j -period ahead “risk-free” bond is priced as

$$Q_{t,j} = \beta^j E_t \left[\frac{\lambda_{t+j}}{\lambda_t} \right] = \beta^j E_t \left[\left(\frac{A_{t+j}}{A_t} \right)^{-\gamma} \right] \quad (7)$$

which implies the j -period bond yield of (8)

$$R_{t,j}^f = \frac{1}{\beta E_t \left[\left(\frac{A_{t+j}}{A_t} \right)^{-\gamma} \right]^{\frac{1}{j}}}$$

Third, consider a j -period-ahead “rental bond” that pays an amount equal to the rent on capital in period $t+j$ only. The price of this bond is equal to:

$$V_{t,j}^R \equiv E_t \beta^j \frac{\lambda_{t+j}}{\lambda_t} R_{t+j} = \beta^j \frac{\alpha H^{1-\alpha}}{A_t^{-\gamma}} E_t A_{t+j}^{1-\gamma} \quad (9)$$

Similarly, we can price a j -period-ahead wage bond that pays an amount equal to the wage in period $t+j$ only.

$$V_{t,j}^W \equiv E_t \beta^j \frac{\lambda_{t+j}}{\lambda_t} W_{t+j} = \beta^j \frac{(1-\alpha)H^{-\alpha}}{A_t^{-\gamma}} E_t A_{t+j}^{1-\gamma} \quad (10)$$

With these expressions, and an arbitrary process for A_t , we can derive closed-form expressions for the asset prices.

Our goal is to find a constant discount rate that can be used to value obligations linked to wages at time $t+j$. To do this, we define the internal rate of return of the wage or rental bonds as the discount rate that gives the wage (or rental) bond an expected present value equal to its price. This is given by the $R_{t,j}^W$ that satisfies the equation

$$V_{t,j}^W = E_t(W_{t+j}) / (R_{t,j}^W)^j$$

In contrast, the actuarial price of the wage bond is given by the expected value of wages at time $t+j$ discounted to the present by the risk-free rate

$$V_{t,j}^{W,A} = E_t(W_{t+j}) / (R_{t,j}^f)^j$$

The key quantity required for computing the market's valuation of wage-linked obligations is the ratio of these prices. This ratio, which we call RA, is given by

$$RA = \frac{V_{t,j}^W}{V_{t,j}^{W,A}} = \left(\frac{R_{t,j}^f}{R_{t,j}^W} \right)^j$$

Solving for this value (it is the same for both rental and wage bonds) yields

$$R_{t,j}^W = \frac{1}{\beta} \left(\frac{E_t(A_{t+j})}{E_t \left[\left(\frac{A_{t+j}}{A_t} \right)^{-\gamma} A_{t+j} \right]} \right)^{\frac{1}{j}} = R_{t,j}^R$$

From this, we can now compute the (annualized) wage bond risk premium:

$$\frac{R_{t,j}^W}{R_{t,j}^f} = \left(\frac{E_t \left[\frac{A_{t+j}}{A_t} \right] E_t \left[\left(\frac{A_{t+j}}{A_t} \right)^{-\gamma} \right]}{E_t \left[\left(\frac{A_{t+j}}{A_t} \right)^{1-\gamma} \right]} \right)^{1/j} \quad (12)$$

Let $\Delta a_t = \log \left(\frac{A_t}{A_{t-1}} \right)$, be the log of productivity growth. Suppose for now that cumulative

productivity growth between periods t and $t+j$ is log-normally distributed, conditional on the information available to agents at time t . That is, assume that

$\sum_{i=1}^j \Delta a_{t+i} \equiv \Delta a_{t,t+j} \sim N \left(E_t \Delta a_{t,t+j}, \sigma_{a,t,t+j}^2 \right)$. Using standard formulas, the log of expression

(12) (the log-risk premium) can be written as

$$r_{t,j}^W - r_{t,j}^f = \frac{1}{j} \left[(1-\gamma) E_t \Delta a_{t,t+j} + \frac{(1+\gamma^2)}{2} \sigma_{a,t,t+j}^2 - (1-\gamma) E_t \Delta a_{t,t+j} - \frac{(1-\gamma)^2}{2} \sigma_{a,t,t+j}^2 \right] = \gamma \frac{\sigma_{a,t,t+j}^2}{j}$$

The ratio of the risk-corrected to the actuarial prices for the wage bond is therefore given by

$$RA = \exp\{j(r_{t,j}^f - r_{t,j}^w)\} = \exp\{-\gamma\sigma_{a,t,t+j}^2\}.$$

For future reference, we also compute the 1-period equity risk premium in the model. The one period return on stock is given by

$$R_{t+1}^e = \frac{D_{t+1} + V_{t+1}}{V_t} = \frac{A_{t+1} \left(1 + E_{t+1} \sum_{j=1}^{\infty} \beta^j \left(\frac{A_{t+1+j}}{A_{t+1}} \right)^{1-\gamma} \right)}{A_t E_t \sum_{j=1}^{\infty} \beta^j \left(\frac{A_{t+j}}{A_t} \right)^{1-\gamma}}$$

The equity premium is given by

$$\frac{E_t R_{t+1}^e}{R_{t+1}^f} = \frac{E_t \left[\frac{A_{t+1}}{A_t} \left(1 + E_{t+1} \sum_{j=1}^{\infty} \beta^j \left(\frac{A_{t+1+j}}{A_{t+1}} \right)^{1-\gamma} \right) \right]}{E_t \sum_{j=1}^{\infty} \beta^j \left(\frac{A_{t+j}}{A_t} \right)^{1-\gamma}} \beta E_t \left[\left(\frac{A_{t+1}}{A_t} \right)^{-\gamma} \right]$$

A “News” Process for A

In this subsection, we describe a “news” process for A, in which agents have information about future productivity that is not captured by the observed history of productivity. Recently, such models have been widely explored in the macroeconomics literature, including prominent studies by Beaudry and Portier (2004), Schmitt-Grohé and Uribe(2010), and Jaimovich and Rebelo (2009).

We suppose that in each period t , agents get “news” about future productivity growth in all subsequent periods. That is, in each period t they observe an infinite sequence of i.i.d shocks $\epsilon_t = \{\epsilon_t^t, \epsilon_t^{t+1}, \epsilon_t^{t+2}, \epsilon_t^{t+3}, \dots\}$. These shocks represent news that arrives in period t about productivity in periods $t, t+1, t+2$, etc. Now, assume that productivity growth in period t is given by the sum of the news shocks regarding period t that are realized in period t and all previous periods. Period t productivity growth is then given by the infinite sum

$$\Delta a_t = g + \sum_{j=0}^{\infty} \epsilon_{t-j}^t \quad (14)$$

where the subscript denotes the period in which the shock is learned by agents, and the superscript denotes the period in which the shock realizes; that is, the period in which it affects productivity growth. The parameter g measure the drift in the technology process, and therefore average wage growth, but will have no impact on the risk-correction implied by the model.

Many assumptions about the importance of news at different horizons are possible, so long as the variance of the combined shocks is finite. For simplicity, and to better match the asset pricing model we present in the next session, we assume that productivity in period t is completely predetermined, but that news shocks for future periods have exponentially decaying variances. The sequence of variances for the shocks ϵ_t is therefore given by $s_\epsilon = \{0, \sigma_\epsilon^2, \rho\sigma_\epsilon^2, \rho^2\sigma_\epsilon^2, \rho^3\sigma_\epsilon^2, \dots\}$.

Because of the simple structure of our model and the process for technology, we can derive a closed-form solutions for all of the relevant variables. Let $\sigma_{a,t+i}^2 = \text{var}_t(\Delta a_{t+i})$ be the variance of the growth rate of productivity between periods $t+i-1$ and $t+i$, conditional on the information available in period t . Then, from the process in (14), we have

$$\begin{aligned}\sigma_{a,t+1}^2 &= E_t[(\epsilon_{t+1}^{t+1})^2] = 0 \\ \sigma_{a,t+2}^2 &= E_t[(\epsilon_{t+1}^{t+2})^2] + E_t[(\epsilon_{t+2}^{t+2})^2] = 0 + \sigma_\epsilon^2 \\ \sigma_{a,t+3}^2 &= E_t[(\epsilon_{t+1}^{t+3})^2] + E_t[(\epsilon_{t+2}^{t+3})^2] + E_t[(\epsilon_{t+3}^{t+3})^2] = 0 + \sigma_\epsilon^2 + \rho\sigma_\epsilon^2 \\ &\dots\end{aligned}\tag{15}$$

$$\sigma_{a,t+i}^2 = \sigma_\epsilon^2 \frac{1 - \rho^{i-1}}{1 - \rho}\tag{16}$$

To compute the risk-premium (and corresponding risk adjustment) we need $\sigma_{a,t,t+j}^2$. But, since all shocks are i.i.d, this is just the sum of the per-period shock variances:

$$\sigma_{a,t,t+j}^2 = \sum_{i=1}^j \sigma_{a,i}^2 = \frac{\sigma_\epsilon^2}{1 - \rho} \sum_{i=1}^j (1 - \rho^{i-1})\tag{17}$$

This implies that the log risk premium on wage bonds is given by:

$$r_{t,j}^W - r_{t,j}^f = \gamma \frac{\sigma_\epsilon^2}{1-\rho} \frac{\sum_{i=1}^j (1-\rho^{i-1})}{j} \quad (18)$$

To compute the risk-adjustment in the news model, we substitute the wage bond risk premium from equation (18) into the expression for RA to get

$$RA(j, \rho, \gamma, \sigma_\epsilon) = \exp\left\{-\gamma \frac{\sigma_\epsilon^2}{1-\rho} \sum_{i=1}^j 1-\rho^{i-1}\right\}$$

Finally, the expression for the equity premium can be simplified using the fact that technology growth between periods t and $t+1$ is known in time t , We get an expression for the equity premium as well

$$\begin{aligned} \frac{E_t R_{t+1}^e}{R_{t+1}^f} &= \beta (\Delta a_{t,t+1})^{1-\gamma} \frac{1 + E_t \sum_{j=1}^{\infty} \beta^j (\Delta a_{t+1,t+1+j})^{1-\gamma}}{E_t \sum_{j=1}^{\infty} \beta^j (\Delta a_{t,t+j})^{1-\gamma}} \\ &= \frac{\beta (\Delta a_{t,t+1})^{1-\gamma} + E_t \sum_{j=1}^{\infty} \beta^{j+1} (\Delta a_{t,t+1+j})^{1-\gamma}}{E_t \sum_{j=1}^{\infty} \beta^j (\Delta a_{t,t+j})^{1-\gamma}} \\ &= \frac{E_t \sum_{j=1}^{\infty} \beta^j (\Delta a_{t,t+j})^{1-\gamma}}{E_t \sum_{j=1}^{\infty} \beta^j (\Delta a_{t,t+j})^{1-\gamma}} \\ &= 1 \end{aligned}$$

The one-period equity premium is nil.

B. A reduced-form model of dividends and wages

As an alternative to our general equilibrium model, we next consider a reduced-form model, which simply assumes a strong positive long-run relationship between average economy-wide wages and dividends (and therefore stock prices), but a short-run correlation that is zero.⁵ Authors, such as Benzoni et al (2007), argue that it is an

⁵ As mentioned above, this model was the basis of related work by Geanakoplos and Zeldes (2010). See also Lucas and Zeldes (2006) for a similar model of wage-stock cointegration.

empirically reasonable description of the link between the stock market and wage, making it a good candidate for the risk-neutral pricing approach established by Cox, Ross, and Rubenstein (1979). Fortunately, as we show, the two approaches can be parameterized to yield identical results.

Log real dividends (d) are assumed to follow a geometric random walk:

$$d_{t+1} - d_t = (g_d - \frac{\sigma_d^2}{2}) + \sigma_d z_{d,t+1} \quad (19)$$

The dividend growth shock, $z_{d,t+1}$, is assumed to be standard normal.⁶ As in Benzoni, et. al., we assume a pricing kernel with a constant price of risk, λ . This implies a constant price-dividend ratio, and therefore a constant dividend yield, δ .⁷ The constant price-dividend ratio implies that stock prices also will follow a geometric random walk that is identical, up to a constant, to the process for dividends. In this model, the total realized real stock return (r^s) equals the dividend yield *plus* the growth in real dividends.⁸

$$r_{t+1}^s = \delta + (d_{t+1} - d_t) = (g_d + \delta - \frac{\sigma_d^2}{2}) + \sigma_d z_{d,t+1} \quad (20)$$

We further assume a process for the log of real wages such that wages and dividends are cointegrated; under our assumptions this also implies that wages and stock prices are cointegrated. We adopt the key equation in Benzoni et al (2007) in which wage growth is assumed to be a function of 1) deterministic wage growth (g_w), 2) the current-period deviation from the long term average wage-dividend ratio ($w_t - d_t - \overline{wd}$) and 3) an i.i.d. wage growth shock (z_{t+1}):

⁶ Equation (1) therefore implies a representation of dividend levels with log-normal shocks and expected growth in the level of dividends equal to g_d .

⁷ We can see this from the present value relationship

$$P_0 = E \sum_{t=0}^{\infty} (1 + r^f + \lambda \sigma_d)^{-t} D_t = D_0 \sum_{t=0}^{\infty} (1 + r^f + \lambda \sigma_d)^{-t} (1 + g_d)^t$$
. Computing the sum, we have

$$\frac{P_0}{D_0} = \frac{1}{1 - (1 + g_d) / (1 + r^f + \lambda \sigma_d)} \approx \frac{1}{r^f + \lambda \sigma_d - g_d}$$
. In continuous time, the last statement is an exact equality.

⁸ One weakness of this model is that it yields the counter-factual result that stock returns and dividend growth have the same volatility.

$$w_{t+1} - w_t = (g_w - \frac{\sigma_w^2}{2}) - \kappa(w_t - d_t - \overline{wd}) + \sigma_w z_{w,t+1} \quad (21)$$

The parameter κ determines the rate at which the wage-dividend ratio “error corrects” deviations from the long-term wage-dividend ratio, \overline{wd} . In general, we allow g_w to be different than g_d , although we will assume they are equal when we calibrate the model.⁹

To price risky wage bonds, we must link the risk in wage-bond returns to a risky asset (stocks) already priced in the market. We use standard derivative pricing methods, and derive an analytic expression for the value of wage bonds.¹⁰ Throughout, we assume that the risk unrelated to the stock market is unpriced.

We begin by deriving the expected level of wages j periods ahead. Then the process for the log deviation of the wage-dividend ratio from its long-run value can be written as

$$w_{t+1} - d_{t+1} - \overline{wd} = (1 - \kappa)(w_t - d_t - \overline{wd}) + (g_w - \frac{\sigma_w^2}{2} - g_d + \frac{\sigma_d^2}{2}) + \sigma_w z_{w,t+1} - \sigma_d z_{d,t+1}$$

Iterating this forward, and solving for the wage level, we find an expression for *realized* level of wages j -periods ahead, conditional on wages at time t .

$$w_{t+j} = w_t + (1 - \kappa)^j (w_t - d_t - \overline{wd}) + \frac{(1 - \kappa) - (1 - \kappa)^j}{\kappa} (g_w - \frac{\sigma_w^2}{2}) - (j - \frac{(1 - \kappa) - (1 - \kappa)^j}{\kappa}) (g_d + \frac{\sigma_d^2}{2}) + \sigma_w \sum_{i=0}^{j-1} (1 - \kappa)^i z_{w,t+j-i} - \sigma_d \sum_{i=0}^{j-1} ((1 - \kappa)^i - 1) z_{d,t+j-i}$$

The price of the wage bond is the expected value of this expression (in levels) discounted by the appropriate rate. Given i.i.d. shocks, it is straightforward to show that

$$E[W_{t+j}] = W_t \exp \left\{ [(1 - \kappa)^j - 1] (w_t - d_t - \overline{wd}) + c_{j,w} (g_w - \frac{\sigma_w^2}{2}) + c_{j,d} (g_d - \frac{\sigma_d^2}{2}) + \frac{\sigma_w^2}{2} c_{j,v} - \frac{\sigma_d}{2} c_{j,\sigma} \right\}$$

⁹ 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.

¹⁰ Lucas and Zeldes (2006) and Geanakoplos and Zeldes (2010) used a Monte Carlo approach to similar problems, whereas here we are able to derive an analytic solution. The two approaches (analytic and Monte Carlo) give identical results.

where

$$\begin{aligned}
c_{j,w} &= \frac{1-(1-\kappa)^j}{\kappa} \\
c_{j,d} &= j - c_{j,w} = \sum_{i=1}^j 1 - (1-\kappa)^{i-1} \\
c_{j,\sigma_w} &= \frac{1-(1-\kappa)^{2j}}{\kappa(2-\kappa)} \\
c_{j,\sigma_d} &= j - 2c_{j,w} + c_{j,v}
\end{aligned}$$

Actuarial versus market value

The actuarial value of a wage bond is simply the expected value of the average wage in period j discounted to the present using the risk-free rate. In our model, therefore, the actuarial (log) price of the wage bond linked to period j is given by

$$\log(\hat{V}_{t,j}^w) = \log(E_t[W_{t+j}]) - (r^f \cdot j).$$

The approach we take to risk correction requires finding probabilities for the model processes such that the returns on the risky asset (stocks) are equal in expectation to risk-free returns, but variances are unchanged. Fortunately, this is easily achieved by replacing the “drift” term in the dividend process with a value \tilde{g}_d , such that

$$\tilde{g}_d + \delta - \frac{1}{2}\sigma_d^2 = r^f, \text{ so that the expected return on stocks equals the risk-free rate of}$$

return. The result is again discounted by the risk-free rate, so that the *risk-adjusted* price for the wage bond is given by $\log(\tilde{V}_{t,j}^w) = \log(\tilde{E}_t[W_{t+j}]) - (r^f \cdot j)$, where \tilde{E} represent expectations taken under the risk-neutral probabilities.

We show in Appendix A that ratio of the risk-corrected to the actuarial wage bond price is given by

$$\widetilde{RA}(j, \kappa, E[r_t^s] - r^f) = \exp\{-c_{j,d}(E[r_t^s] - r^f)\} \quad (22)$$

Equation (22) says that the ratio of market to actuarial wage bond prices is a linear function of the risk premium, with slope-coefficient $-c_{j,d}$ that depends on the strength of cointegration between wages and stocks. It is natural that wage growth

does not appear in this expression, since the deterministic portion of wage growth is unaffected by the shift to risk-neutral probabilities for stocks. Equation (4) captures the full risk-adjustment implied by our model.

C. Estimation

Several papers have estimated the correlation between stock prices or returns and wages or wage growth. We reiterate that a finding of small contemporaneous correlations commonly obtained in papers using factor pricing methods may not convey the full story because correlations may be small in the short run but large in the long run. Jermann (1999) essentially finds no correlation between labor income growth and stock returns at short horizons, but finds a strong correlation for horizons greater than 20 years. For 6-year periods and real per-full-time-equivalent worker wage growth, Krueger and Kubler (2006) find a correlation of -0.38 using a CPI-deflated NYSE/AMEX value-weighted portfolio, and Krueger and Kubler (2002) find a correlation of 0.418 when using the S&P500 and the NIPA deflator for total consumption expenditures. Using historical data for the United Kingdom, Khorasanee (2009) obtains a negative contemporaneous correlation between equity returns and wage growth but a positive correlation between .26 and .4 when equity returns are lagged by one to three years.

If wages and dividends are cointegrated then they may exhibit little correlation in the short run, but the difference in their levels will predict the growth rate of one or both. To test for cointegration, Benzoni et al (2007) apply an augmented Dickey-Fuller test to the demeaned log difference between the dividends and per capita compensation. Looking solely at the post-World War II period, they cannot reject the hypothesis of a unit root in the process for this difference, but rejection is possible when including all years from 1929 to 2004 and particularly strong when including lagged changes in the difference to allow for serial correlation of the error term. The estimated coefficient on the lagged difference in these same regressions yields κ , the most important parameter choice from the perspective of this paper. They estimate κ to be between .05 and .2 and take 0.15 as their baseline value.

We follow an estimation strategy similar to that of Benzoni et. al. (2007). Specifically, we construct monthly dividends from the Value-Weighted Return and Total

Market Value series of the Center for Research in Security Prices, annualize these to use in conjunction with the annual per-capita compensation data from the National Income and Product Accounts (NIPA) from the Bureau of Economic Analysis, and deflate both using the NIPA Personal Consumption Expenditures Deflator. For the denominator of per-capita compensation (which corresponds to wages in the model) we try both total population, which replicates the results in Benzoni et. al. (2007), and employment. Using employment we get similar results from basic specifications, and in regressions including lagged differences the use of employment increases the adjusted R^2 from .22 to .31 and gives values of κ above .3. Employment data is only available from 1939 onward, so we also try using total population as the denominator over this period and obtain estimates similar to those of Benzoni et. al. (2007) but with larger adjusted R^2 . We therefore maintain .15 as the baseline value for κ and then study robustness to other values.

D. Calibration

With analytical expressions for each model, we can now compare the risk adjustments from the two models.

$$RA(j, \rho, \gamma, \sigma_\epsilon) = \exp\left\{-\gamma \frac{\sigma_\epsilon^2}{1-\rho} \sum_{i=1}^j 1 - \rho^{i-1}\right\} \quad (23)$$

$$\widetilde{RA}(j, \kappa, E[r_t^s] - r^f) = \exp\left\{(r^f - E[r_t^s]) \left(\sum_{i=1}^j 1 - (1-\kappa)^{i-1}\right)\right\}$$

Inspecting the two formulae, we see that the summation terms in each expression are equivalent if $\rho = 1 - \kappa$. Picking ρ and κ to satisfy this restriction, we can equitize the two risk adjustments by selecting γ and σ_ϵ so that the term $\gamma \frac{\sigma_\epsilon^2}{1-\rho}$ corresponds to the equity premium in the asset pricing model. In our computations, we take the length of a period to be one-quarter, although all parameters values are expressed in annualized rates.

Computing the risk correction from the risk-neutral pricing model requires

calibration of only two parameters: excess stock returns and the value of the cointegration parameter, κ . Using monthly S&P data from 1959 through 2008, we compute an average annual equity premium (over three-month Treasury bills) of approximately 5 percent. As described in the preceding section we take .15 as our baseline value for κ .

As demonstrated above, the GE news model can be calibrated so that the corresponding risk correction is exactly equal to that derived from the reduced-form pricing model. In the pricing model, the equity premium (parameterized at 5 percent) determines the overall level of the risk correction, while κ determines how the correction varies for wage bonds of different horizons. Since the parameter ρ in the expression for news model correction corresponds exactly to $1 - \kappa$ in the pricing model expression, we choose $\rho = 1 - \kappa$.

When $\rho = 1 - \kappa$, setting $\gamma \frac{\sigma_\epsilon^2}{1 - \rho} = E[r_t^s] - r^f$ causes the risk-adjustment in the two models to be identical. In order to identify particular values for the news variance and risk aversion, we calibrate the model so that the unconditional volatility of output growth in the model corresponds to that in the data. Recall that in the news model, output and technology move one-for-one in percentage terms, so that $\sigma_A^2 = \sigma_Y^2 = \frac{\sigma_\epsilon^2}{1 - \rho}$. Using data from 1947 to 2008, we compute the standard deviation of real per-capita output growth to be 2.8 percent. From this we compute a value for $\gamma = \frac{.05}{.028^2} \approx 63.8$. Finally, we find that $\sigma_\epsilon = (1 - \rho)\sigma_Y = .15 \times .028 = .0042$.

The implied value for γ above implies a degree of risk aversion that is greater than is generally thought reasonable. This difficulty corresponds to the well-known “equity premium puzzle” in economics: it is very hard to generate large equity premia in models with standard preferences and reasonable volatility in consumption.

Although they imply the same risk-correction, these two models do not have identical implications for all prices. In particular, the news model implies a time-varying risk-free rate, which is assumed to be constant in the risk-neutral pricing model. Excessive volatility in the risk-free rate is a second well-recognized asset pricing puzzle

generated by other basic models. Fortunately we need not choose between these alternatives because our goal is to price wage-indexed securities, and either approach leads to the same conclusions, given our assumptions about productivity.

E. Wage bond prices: results

Figure 1 plots the term structure of risk premia that should be built in to the discount rate for wage-linked claims, i.e. the difference between $\bar{r}_{t,j}^W$ rate and the risk-free rate under market value and actuarial approaches. The actuarial approach simply discounts using the risk-free rate in all periods, with no additional adjustment for risk.

Figure 1 reveals two important properties of the term structure of risk premia. First, the risk premia are always positive, i.e. the appropriate discount rate is always higher than the risk-free rate. Second, the risk premia are increasing with the horizon, i.e. the term structure is upward sloping. This is because cointegration implies an increasingly strong correlation between current stocks and future wages. As a result, the risk premium built in to the discount rate goes from zero for wage bonds maturing in the very near future to the equity premium on stocks for wage bonds maturing in the distant future. This implies that correctly timing the resolution of wage uncertainty associated with future Social Security cash flows will be crucial to determining the proper risk-adjustment.

Figure 1b shows the sensitivity of these results to the choice of equity premium. A lower assumed equity premium results in a lower term structure of risk-premia on wage bonds.

Figure 2 plots the ratio $RA(j) = \widetilde{RA}(j)$ under the parameterization described above. The risk-correction for wage bonds is increasing in j . There is no adjustment to the price of wage bonds maturing in just one period. For a horizon of 20 years, the adjustment is about 50%; and for wage bonds maturing at the end of a 75-years, risk adjustment reduces prices by over 96%.

F. Other Aggregate Risks / Pricing Issues

In this section, TBA, we will address the issues of 1) inflation risk, 2) demographic risk, and 3) questions about marginal prices being applied to large quantities.

G. Comparison to factor analysis approach

If our models that imply wage-stock market cointegration are true, then care must be exercised when applying the standard factor analysis approach to pricing wage bonds (as in Blocker et al., 2009 and Goetzmann, 2008.) In the factor-analysis approach, the price of an asset (and thus the proper risk-correction) is determined by regressing the asset's returns on a set of "risk-factor" portfolios and, in essence, using the corresponding "betas" to determine the proper discount rate. In pricing wage bonds, it is tempting to regress the growth in wages on contemporaneous (or nearly contemporaneous) stock returns. In general, however, this will give misleading results because stock market returns and wage bond returns may both contain information about future wage growth beyond what is contained in contemporaneous wage growth.

Using the reduced-form model, which assumes only one "market" risk factor, we can derive the asymptotic regression coefficients for both the correct factor pricing regression (wage bond returns on stock returns) and the potentially misleading regression (wage growth on stock returns.) Consider first estimating the "correct" relations between wage bond returns and stock returns, $\Delta \log(\hat{V}_{t,j}^W) = \beta_j r_t^s + \varepsilon_t^j$ for the wage bond linked to wages in year j . The cointegration of wages and stocks suggests that wage bond prices should increase in response to an increase in the stock price. Furthermore, at long horizons, the change in stock prices will be fully reflected in wage levels (long-run correlation is one), implying a larger coefficient for bonds linked to far-off wage levels. In fact, in Appendix B, we establish that for $h=1$ the coefficient of this regression is $1 - (1 - \kappa)^{(j-1)}$.

Now, consider estimating the contemporaneous relation $\Delta w_t = \varphi r_t^s + v_t$. Because stock returns are i.i.d. and wage growth from period t to $t+1$ depends only on the deviation of the period t wage-dividend level, the regression coefficient tends asymptotically to zero. Under our model, naively substituting wage growth for the wage

bond return would give no risk correction at all!¹¹ Estimating the relation at one lag, $\Delta w_{t+1} = \phi r_t^s + v_{t+1}$ recovers the coefficient of cointegration, κ , asymptotically, but the implied risk-correction will still vastly understate the proper value for bonds at any horizon greater than 1.

Figure 3 shows the implied discount rates under all three approaches to risk-correction. Factor analysis allows for a risk-correction that makes the implied discount rate greater than the risk-free rate. However, Blocker et. al. (2008) estimate market risk from short-term correlations and assume this same adjustment applies at all horizons. This is equivalent to applying the one-year wage bond discount rate to wage bonds of all maturities, which fails to account for correlation that increases with maturity. Again, wage bonds with distant maturities are riskier, and therefore should be discounted at a higher rate.

The key issue here is that of predictability. If wage bond prices were a constant multiple of the wage level, then the contemporaneous growth regression would be appropriate. In any model where the current wage level is not a sufficient statistic for the expectation of future wages, however, this regression will be misleading. In our model, the current stock price predicts future wage growth, even after controlling for the wage level, and this causes wage bond returns to be quite different than wage growth. Substituting growth rates for returns may be plausible under other models, but clearly gives misleading results in our environment.

In related work, Koehler and Kotlikoff (2009) advocate computing the market value of Social Security flows. Their approach is similar to that of Blocker et. al. (2008), but also addresses some of the concerns raised by Geanakoplos and Zeldes (2010). For one set of calculations, they run a separate regression to estimate risk at each horizon, which potentially corresponds to the horizon specific risk adjustment we compute using our model. However, data limitations make it difficult to interpret their results. First, their dataset has only 22 years of observations, requiring them to extrapolate to longer horizons. They also use tax and benefit payments as their dependent variables, rather than the wage growth that determines the levels of these

¹¹ This is an asymptotic result because this regression is biased in finite sample. Using data simulated from our model, we generate an average coefficient (and downward bias) of -.035 from 100,000, 50-year simulations.

payments, which is likely to produce spurious correlations induced by historical changes in program parameters like payroll tax rates and benefit formulas. Their procedure leads to risk adjustments that *increase* the value of both taxes and benefits well above the actuarial values reported by SSA, suggesting a negative correlation between market securities and future program flows. Because of the limitations of the data, we find it difficult to determine if these correlations are artifacts of the historical experience (including policy changes) or a true correlation that can be extrapolated going forward. We continue to believe that the weight of the evidence (empirical and theoretical) favors a positive long-run correlation between real wages and risky assets, and therefore that risk-adjustment should lead to a higher, not a lower, discount rate.

III. Constructing market value measures of Social Security

In this section, we use the derived prices for the wage bonds to compute the market value of Social Security's assets and obligations.

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 \$106,800 in 2010). 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 (2009) 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 60.¹² 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. Defining measures of unfunded obligations

We start by examining the three most common measures of the actuarial status of Social Security: the Immediate Transition Cost, the Closed Group Transition Cost, and the Open Group Unfunded Obligation. The OACT publishes a recurring actuarial note (Wade et. al. 2009) that presents these measures for various horizons. The published measures are for the entire OASDI program, which includes spouses, survivors, and disability claimants. To simplify our analysis (and make the risk adjustment as clear as possible), we focus instead on the OAI portion of benefits (i.e. we exclude disability, spouse, and survivor payments) and treat all benefits as being paid to retired workers. We scale these benefits up to correspond with the OACT’s OASI tax revenue estimates (that include a portion related to survivors that we cannot back out), but we refer to our measures as OAI measures because we perform risk adjustments according to the distribution of OAI benefits. In addition to the 75-year horizon the OACT uses for the Open Group Unfunded Obligation, we also calculate the value of this measure for a 500-year horizon to approximate the imbalance for the infinite horizon. We also present a new Future Transition Cost measure that includes all cash flows for the 75-year Open Group measure plus all benefits beyond 75 years that will have been accrued as of 75 years from now. We view this as a more appropriate measure: unlike the 75-year Open Group measure, the FTC measure pays all accrued benefits associated with collected taxes. Note that the FTC could easily be calculated for any horizon to determine the cost of running the program until that date and then shutting down. Details of the construction of these measures can be found in Appendix

¹² 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.

C.

The Immediate 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 the status measures used for private pensions. As of January 2009 the SSA actuarial value of this shortfall stood at \$18.9 trillion for the entire OASDI program.

The Closed Group Unfunded Transition Cost is a 100-year measure of the difference between all benefits that will be paid to individuals that are already of working age and the taxes that will be paid by these workers. The closed group measure differs from the Immediate Transition Cost because it includes the future tax payments and future 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 total unfunded obligation over the next 100 years is \$16.3 trillion.

The Open Group Unfunded Obligation considers all prospective flows, including taxes and benefits for those not yet in the workforce, and is reported for both a 75-year horizon and the infinite horizon. For the 75-year period the open group 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 75-year OASDI open group obligation currently stands at \$5.3 trillion, while the OASDI unfunded obligation for the infinite horizon is estimated at \$15.1 trillion.¹³

Figure 4 depicts projected real cash flows for our OAI version of each of the three published measures. Yellow bars represent benefits that have accrued to-date and are thus included in the Immediate Transition Cost. Green bars give the cash flows from current participants included in the Closed Group measure, but not in the ITC. Finally, blue bars represent cash flows from future participants, which are added to the Closed Group measure to derive the Open Group measure. Because workers pay taxes early in life and receive benefits later in life, the majority of benefit payments over the horizon

¹³ For more details on these measures, see Goss (1999) and Wade, Schultz, and Goss (2009).

in the figure go to current workers, while the majority of tax payments come from future workers. This difference will be critical to the size and the direction of risk-adjustment.

Figure 5 presents the same cash flows, but with the values representing actuarial present values rather than constant dollars. This figure, therefore, represents the *actuarial perspective* on the relative value of the various social security cash flows. *Lastly*, Figure 6 extends Figure 5 to show the first 100 years of the Open Group cash flows for our infinite horizon projection. After the initial 75-year period for which we have data we project that real tax revenues continue to grow at 1.48% annually and real benefits grow by 1.87% annually, as they do over the last 5 years of the 75-year horizon.

C. Translating cash flows into wage bonds

To price these cash flows using the results from the pricing models, we need to construct the portfolio of wage bonds that will match the projected cash flows of the system. Because Social Security taxes are essentially a flat percentage of current income, the government's long position in wage bonds (i.e. the number of wage bonds it owns) with maturity j is equal to the taxes to be paid j years in the future measured in units of period j wages. The value as of today of period j taxes will then equal the number of period j wage bonds held multiplied by the price today of a period- j wage bond.

Valuing benefits is slightly more complicated because the risk-adjustment applied to benefit payments in a given year depends on which cohort receives them. Specifically, since benefits paid to a particular cohort are proportional to the average wage when that cohort reaches age 60 (excluding the special benefits mentioned above), the government's short position in wage bonds maturing in year j is equal to the annuity value of benefits paid to the cohort turning 60 j periods from today. Appendix C describes how we allocate projected annual cash flows across the different cohorts.

Figure 7 depicts the actuarial present value of gross wage bond assets and liabilities by year of maturity (and by this we mean the year in which uncertainty about the cash flow is resolved). Because future tax flows are linked to wages in the same year, the pattern of long positions (positive bars) is unchanged from Figure 5. On the other hand, benefit payments typically occur long after their associated wage-risk is

resolved. Thus, represented in terms of their wage-bond maturity date, benefit cash flows are shifted significantly towards the present. Indeed, the amount of riskless “0-year” wage bonds (corresponding to the benefits of individuals who reached age 60 by 2008) is by far the largest position held by the government.

Net wage bond holdings (taxes less benefits) are shown in Figure 8 for the published measures of the system’s financial status. For both the Immediate Transition Cost and the Closed Group these net holdings are negative at all horizons. For the 75-Year Open Group, which includes taxes from future workers whose benefits lie outside the 75-year window, net holdings of wage bonds with long maturities are positive. Our estimate of Social Security’s actuarial deficits is computed by (equivalently) summing up the present values in either Figure 5 or Figure 7 and subtracting the current \$2.4 trillion value of OASDI Trust Fund.

E. Defining market value measures of financial status

Our goal in this paper is to emphasize the importance of risk-correction, holding constant other aspects of the current actuarial procedure. We showed above that the ratio of market to actuarial wage bond prices is a function only of the time-horizon and the risk-premium earned by stocks. To obtain the market value of each of the three measures of unfunded obligations we need only multiply that measure’s net holdings of wage bonds maturing in each year $t+j$ by the ratio of market to actuarial prices for j -year wage bonds. The sum of these risk-adjusted wage bond holdings (less the value of the Trust Fund) is the market value of the unfunded obligation. Incorporating the risk-adjustment suggested by our model into the traditional actuarial approach is remarkably simple and only requires estimates of κ and the (real) equity risk-premium.¹⁴

IV. Estimates of the market value of Social Security

Here we examine the effects of risk adjustment on the four measures of financial status defined above.

A. Immediate Transition Cost

¹⁴ We set h equal to 1/100th of a year, but in practice the choice of h has minimal effects. Setting h as large as a full year changes none of our results by more than .5%.

Figure 9a compares the market and the actuarial valuations of wage bonds included in the Immediate Transition Cost. Note that the risk-adjustment reduces the value of the liability for all cohorts under age 60. Since this measure includes no taxes, the gross and net positions are the same and always negative. In this case, risk-correction can only work to reduce the shortfall. Table 1 presents the summary results. Properly accounting for risk reduces the amount of the Immediate Transition Cost by \$2.6 trillion, an 18% reduction.

B. Closed Group

Figure 9b shows the market and actuarial valuations of Closed Group wage bond holdings. The adjustment is even more important for the Closed Group Transition Cost than it is for the Immediate Transition Cost. Table 1 indicates that the risk adjustment dramatically reduces our estimate of the closed group deficit, from \$13.7 trillion to \$8.1 trillion, a 41% reduction.

Why is this effect so strong? To see why, consider any currently-working cohort (recall that benefits are already fixed for retired workers). 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.

Another way of seeing why risk adjustment is particularly “beneficial” for the Closed Group measure is to compare the timing of its flows to those of the Immediate Transition Cost. The value of wage bonds under the Maximum Transition Cost measure is most negative at the shortest maturities corresponding to older workers who have already accrued the majority of their benefits. Holdings become less negative as maturity increases. The Closed Group, though, is short the most wage bonds in maturities around 35-45 years, when the youngest workers of today will finish paying taxes and begin collecting benefits. Risk correction is quite strong for such maturities, with market-to-actuarial ratios of 25% and lower. The total adjustment to the Closed Group measure is therefore quite large.

C. Open Group

Figure 9c shows the value of wage bonds included in the 75-year Open Group measure. In contrast to the other measures, in this case the risk-adjustment gives a larger shortfall than the corresponding actuarial measure. The reason is easy to see in the figure. This measure includes as an asset the taxes of many future workers who receive no benefits within the 75 year window of the Open Group measure. Market valuation discounts these late-maturing wage bonds most heavily, eliminating the vast majority of their value. In contrast, the negative net holdings of wage bonds with earlier maturities are discounted less heavily, making the overall total more negative (i.e. showing a larger deficit) under market valuation.

Figure 9d shows the first 150 years of our infinite horizon measure. Unlike the 75-year measure, which includes tax payments of future generations with little or no offsetting benefits paid to these cohorts, the more forward-looking measure contains wage bonds with negative net actuarial values at all horizons. As a result, the risk adjustment shrinks these debts towards zero and makes the total unfunded obligation less negative. Table 1 also shows the size and direction of the adjustment for Open Group measures of three different lengths. As discussed previously, the risk adjustment increases the size of the 75-year Open Group Unfunded Obligation by 40%. In contrast, a 100-year measure is shrunk to 92% of its actuarial value, and a 500-year measure (which includes over 99% of tax and benefit flows for the infinite horizon) to just 32% of its nearly \$20 trillion actuarial value. Furthermore, the first two columns of Table 1 show that cash flows beyond the 100-year horizon add almost nothing to the market value of either tax revenues or benefit obligations because of the strong effect of the compounded risk correction.

Understanding the impact of the risk adjustment requires distinguishing between two competing effects. The first effect, which tends to increase the measured shortfall, can be seen by considering the risk adjustment to cash flows in a particular future calendar year $t+j$. All benefits (outflows) being paid that year will be determined prior to $t+j$, and on average, much sooner. In contrast, all cash inflows come from labor income taxes, which depend on the wage level realized in period $t+j$ exactly. Since the tax revenues are tied to wages further in the future, they are more risky and are therefore subject to a larger risk correction. This effect always works to increase the measured

shortfall in the system.

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. Therefore, unlike the first effect, the direction of this effect depends on the starting point. If the system were in surplus, this effect of risk adjustment would shrink the surplus. However, since our system is currently in deficit, this second effect shrinks the deficit and thus works in the opposite direction of the first effect. For the 75-year measure, the first effect dominates, causing risk-adjustment to increase the magnitude of the shortfall, whereas for the infinite horizon measure the second effect dominates, causing the shortfall to decrease.

D. Future Transition Cost

Figure 9e shows the value of wage bonds included in the Future Transition Cost measure. We compute this new measure as an additional alternative to the standard Open Group measure that contemplates paying all accrued obligations at the end of the 75-year period rather than taxing future workers and giving them little or no benefits. As with all other measures besides the 75-year Open Group, the risk-adjusted value of the Future Transition Cost gives a smaller shortfall than the corresponding actuarial measure. The actuarial value of this measure is \$13.1 trillion. At \$6.4 trillion, the market value of the measure is 49% of actuarial value. Accounting for benefits accrued by the end of the period eliminates the large positive net holdings of the 75-year Open Group Unfunded Obligation and causes the risk correction for the Future Transition Cost to work in the same direction as for the other measures.

V. Policy Implications

In this section, we examine two policy implications of our findings about risk adjustment.

A. How large an adjustment is necessary to balance Social Security?

The OACT publishes each imbalance measure in level form and also as a percentage of projected taxable payroll and GDP. For the full OASDI program as of January 1, 2009 the 75-year Open Group Unfunded Obligation was 1.9% of 75-year payroll and 0.7% of GDP over that period, the infinite horizon Open Group Unfunded Obligation was

3.4% of all future payroll and 1.2% of all future GDP, the Closed Group Transition Cost was 5.8% of payroll and 2.1% of GDP, both over the next 75 years, and the Maximum Transition Cost was 6.7% of payroll and 2.4% of GDP over the next 75 years. These ratios facilitate year-to-year comparisons and indicate the level of additional taxation that would be needed to rebalance the system if behavioral responses were small. In this section we compare the percentage increases in taxes necessary to balance the system under the actuarial and market value approaches. We will also consider the required percentage decreases in benefits.

To determine the necessary payroll tax increase we divide the amount of the actuarial imbalance for a particular measure by the present value of the relevant projected payroll for that measure.¹⁵ The market value of each measure's payroll is simply the actuarial value scaled down using the markdown series derived from our preferred parameter values.

Table 3 presents the changes required to rebalance the system according to each measure. The first column of numbers shows the actuarial and market values of the imbalance measures reported in the previous tables. The second column of numbers shows the unfunded obligations as a percentage of the present value of all future payroll based on the OASI payroll tax flows and a payroll tax rate of 12.4%. Based on actuarial estimates it would require a permanent increase in payroll taxes of around 4 percentage points to balance the Closed Group and Transition Cost measures, 1.2 percentage points to balance the 75-Year Open Group measure, and 5.6 percentage points to balance the 500-year Open Group measure. Although the market value of all but the 75-year Open Group measures is less negative than the actuarial value, we find that adjusting for wage risk requires a larger payroll tax under all the measures. The reason is that the additional revenue raised by the tax increase is itself subject to wage risk and hence worth significantly less when subject to the market risk correction.

The next column of Table 3 presents the results of a slightly different percent-of-payroll calculation, using only the payroll relevant for that imbalance measure. This

¹⁵ To obtain covered payroll over the next 75 years we divide projected payroll taxes for each measure by .124, the current payroll tax rate. Beyond the 75-year window we assume the present value of taxes (and hence payroll) changes at the annual rate of .986 as it does over each of the last five years in the data.

obviously makes no difference for the 500-year Open Group measure, but for the Closed Group it implies raising taxes only on the cohorts who remain part of the system, and for the Immediate Transition Cost this measure cannot be calculated because there is no payroll to tax. Using actuarial value, the switch from 500-year payroll to relevant payroll increases the required tax hike by about 50% for the 75-Year Open Group and Future Transition Cost, and for the Closed Group it would be necessary to increase the tax by 12.3 percentage points, effectively doubling the payroll tax. Using market value, on the other hand, putting relevant payroll in the denominator does not increase the necessary tax by very much relative to the 500-year payroll calculation because the market value of uncertain revenue received in more than 75 years is negligible. The Closed Group measure does increase from 8.6% to 11.2% because so much of the Closed Group tax revenue is received in the very near future, but note that this 11.2% is smaller than the 12.3% actuarial value for this measure because the market adjustment shrinks the total imbalance by more than it shrinks this near-term tax revenue.

The final column of Table 3 displays the percentage reduction in benefits that would balance the system according to each measure. To determine the necessary benefit reduction for each measure we find x such that for that measure $(1-x) \times (\text{present value of benefits}) = (\text{present value of taxes}) + (\text{trust fund})$. Under the actuarial approach the required benefit cut is 85.5% for the Immediate Transition Cost, which has no offsetting revenue and only the assets of the Trust Fund, which are about 15% of the shortfall. Continuing with the actuarial approach, benefits must be cut 45.7% for the Closed Group, 12.4% for the 75-year Open Group, and about 30% for the 500-year Open Group and Future Transition Cost. For the Closed Group and Open Group the market adjustment works in the same direction as for the total amount of the imbalance, with larger deficits requiring larger benefit cuts to balance. On the other hand, while the total imbalance under the 500-year Open Group and Future Transaction Cost measures looks better using market value, in both cases market valuation calls for a slightly larger benefit cut than actuarial value. These long-term measures both anticipate paying a considerable amount of benefits after 75 years, as can be seen by comparing the present values of their benefits in the earlier tables to that of the 75-Open Group measure. Market valuation takes a heavy toll on the value of benefits for these long-term measures, whereas benefits under the other measures eventually decline with the

number of surviving cohorts, so the risk adjustment reduces benefits by a smaller percentage.

Market valuation has fairly straightforward effects on actuarial (im)balance measures, generally making them closer to zero, but the policy implications for balancing the system can be more complicated. If taxes are increased then the additional revenue will be subject to risk adjustment, and if benefits are decreased the same risk adjustment can cause this reduction to have a less-than-proportional effect on present value. Whether policy reforms must be more or less severe under market valuation depends on the timing of cash flows to be affected and hence on the choice of imbalance measure.

B. Would shifting from wage-indexing to price-indexing save costs?

One policy question that often arises is whether benefits should be linked to wages (as they are through age 60 under the current program) or instead linked to the overall price index. Some authors have suggested that the cost of providing retirement benefits could be reduced by price indexing (linking future benefits to future inflation) rather than wage-indexing (linking future benefits to future aggregate wages). From the actuarial perspective, this is true: real wages are expected to grow by about 1.1% per year, meaning that, at 30 years, one dollar in benefits today would be expected to grow to equal about $1.01^{30} = 1.39$ times what the benefit would be if it were only inflation indexed.

Our analysis reaches the opposite conclusion: at market prices, under our parameterization, wage-indexed benefits cost substantially *less* than inflation-indexed benefits. Consider a dollar wage- or price-indexed benefit earned today. As described above, in real terms, these streams correspond to a wage and risk-free bond, respectively. Figure 10 compares the present value of wage- and price-indexed cash flows at different horizons. In the very near term, there is indeed a higher price for a wage-indexed cash flow than for a price-indexed cash flow. However, by the fourth year, the riskiness of wages means that they are discounted at a rate that more than compensates for their higher expected value. At longer horizons, the wedge between the two grows larger so that, at the 30 year horizon, a wage-indexed benefit has a market value that is 59% *less* than that of the same price-indexed benefit, in contrast to

the actuarial approach that indicates that the wage-indexed benefit costs 39% more. In terms of market prices, the risk-correction overwhelms the higher expected growth rate of wages.

The intuition for this result is as follows. The risk correction of wage bonds accounts for the (small) probability that at some future date, wages (and productivity) will be extremely low. In this state, paying an inflation-indexed bond will be quite painful, whereas the wage-linked obligation will be much smaller.¹⁶ The conclusions here are the same in both of the models that we consider. However, they depend crucially on assumptions about real wage growth and the equity premium. If real wages were assumed to grow at a faster rate on average, or if the equity premium were less than we assumed, then the costs of wage indexing would be increased relative to price indexing.

Our result above depends crucially on the existence of realizations of the economy in which real wages many periods in the future are lower than they are today. To shed light on how likely these realizations might be in practice, we examine data on real wage growth in the U.S. and a number of other countries.

Evidence shows that U.S. workers have avoided the misfortune of sustained real wage declines, at least over the past 70 years. NIPA real per-employee compensation has declined in 3 of the 66 5-year periods in the data, but not in any 10-year periods, and the Social Security Administration's Average Wage Index has had no 5-year periods of real decline since 1951. However, the U.S. appears to be the exception in this regard. We perform similar calculations for 29 countries using data from the OECD.StatExtracts provided by the Organization for Economic Co-operation and Development. This includes compensation of employees, total labor force, and a consumer price index for each country, allowing us to calculate real per-employee wage growth.¹⁷ Our findings on negative real wage growth are summarized by country and observation in Table 4. The results for the U.S. are comparable to estimates described

¹⁶ As noted above, we have assumed there is no inflation risk in the economy. In practice, the true price of inflation risk is hard measure, with even its sign in some dispute. It seems unlikely that adjusting for inflation risk would change the basic result that wage-indexed benefits are less expensive than price-indexed benefits.

¹⁷ The documentation indicates that "In all cases a lot of effort has been made to ensure that the data are internationally comparable across all countries presented and that all the subjects have good historical time-series' data to aid with analysis."

above, with no long periods of real wage decline. However, for the entire set of countries, we find that fully 10% of the observed 20-year wage growth rates are negative. Germany, Mexico, the Netherlands, and New Zealand all experienced periods of real wage decline lasting 30 years or more. These results suggest that long-term declines in real wages are a possibility, providing some additional credence to our result that wage indexing is less expensive than price indexing.

VI. Conclusions

We argue that market value is the appropriate way to measure both assets and liabilities of the Social Security system, and we estimate the appropriate rates for discounting future risky Social Security flows. We show that the proper discount rate is not the risk-free rate, as used by the Social Security actuaries. Nor is it based on the expected returns of the financial assets held by the plan, as has been common actuarial practice with state and local pensions.¹⁸ Instead, the appropriate rate for discounting cash flows depends on the risk embedded in those cash flows. We show that implementing the appropriate risk-adjustment can be done in a relatively straightforward way, conditional on agreeing on an underlying economic specification of the long-run risks of wages.

Adjusting for risk alters all of the SSA measures of imbalance, reducing the deficit as measured by the Immediate Transition Cost, Closed Group Unfunded Obligation, and the infinite horizon Open Group Unfunded Obligation by between 18% and 68%.

We study three policy proposals: increasing payroll taxes, cutting benefits, and indexing benefits to prices rather than wages. For the first proposal, we find that, despite the smaller measured shortfall, the tax necessary to bring the system back into balance is actually greater when measured with proper risk adjustment. For the second

¹⁸ Among pension actuaries dealing with public pension plans, common practice has been to select a discount rate based on the expected return of the plan's assets. For example, based on GASB recommendation, state and local pension funds typically invest 60-70% of assets in stocks and use a discount rate for liabilities that reflects the higher expected return on this portfolio of assets. (The discount rates used are fairly tightly clustered around 8 percent; see Lucas and Zeldes, 2009, and Novy-Marks and Rauh, 2010.) The Social Security Trust Fund is invested in riskless Treasury bonds. It is therefore difficult to tell whether the SSA's use of a riskless rate to discount cash flows is based on the view that cash flows should be always be discounted using a risk-free rate or whether it is based on the view that cash flows should be discounted using the expected rate of return on the assets in the fund.

proposal, we find mixed results. For the third proposal, we find that on a risk-adjusted basis linking benefits to prices (instead of wages) would actually *increase* the cost of Social Security benefits and thus *increase* the shortfall.

Our results highlight the importance of risk-adjustment for government calculations. In future work, we hope to expand the analysis to examine overall government revenues and expenditures.

APPENDIX

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

A. The ratio of market-to-actuarial, analytical solution

\tilde{P}_t^j is the market price of a wage bond that matures in period $t + j$, and

P_t^j is the actuarial price of a wage bond that matures in period $t + j$.

$$\tilde{P}_t^j = E \left[\frac{\tilde{W}_{t+j}}{W_t} \right] \Big/ (1+r^f)^j \quad \text{and} \quad P_t^j = E \left[\frac{W_{t+j}}{W_t} \right] \Big/ (1+r^f)^j.$$

$$E[W_{t+j}] = W_t \exp \left\{ [(1-\kappa)^j - 1](w_t - d_t - \overline{wd}) + c_{j,w} \left(g_w - \frac{\sigma_w^2}{2} \right) + c_{j,d} \left(g_d - \frac{\sigma_d^2}{2} \right) + \frac{\sigma_w^2}{2} c_{j,v} - \frac{\sigma_d}{2} c_{j,\sigma} \right\},$$

so

$$\log(\tilde{P}_t^j) - \log(P_t^j) = c_{j,w} (\tilde{g}_w - g_w) + c_{j,d} (\tilde{g}_d - g_d).$$

We further assume that $\tilde{g}_w - g_w = 0$, so $c_{j,w} (\tilde{g}_w - g_w) = 0$.

We calibrate the model so that $g_d = E[r^s] - \delta + \frac{\sigma_d^2}{2}$. Similarly, for risk corrected,

$$\tilde{g}_d = r^f - \delta + \frac{\sigma_d^2}{2}.$$

So $\tilde{g}_d - g_d = r^f - E[r^s]$ and the risk correction ratio is

$$\log(\tilde{P}_t^j) - \log(P_t^j) = c_{j,d} (r^f - E[r^s]) = -c_{j,d} (E[r^s] - r^f).$$

In levels $\frac{\tilde{P}_t^j}{P_t^j} = \exp \left\{ -c_{j,d} (E[r^s] - r^f) \right\}$.

B. Comparison to factor analysis

Regression 1: Δw_{t+1} on Δd_{t+1}

$$\text{cov}(\Delta w_{t+1}, \Delta d_{t+1}) = \text{cov} \left(\left(g_w - \frac{\sigma_w^2}{2} \right) - \kappa (w_t - d_t - \overline{wd}) + \sigma_w z_w^{t+1}, \left(g_d - \frac{\sigma_d^2}{2} \right) + \sigma_d z_d^{t+1} \right).$$

g_w, g_d, σ_w and σ_d are constant. w_t, d_t and z_w^{t+1} are orthogonal to z_d^{t+1} , so $\beta = 0$.

Regression 2: $\Delta \ln w_{t+1}$ on Δd_t

$$\text{cov} \left(\left(g_w - \frac{\sigma_w^2}{2} \right) - \kappa (w_t - d_t - \overline{wd}) + \sigma_w z_w^{t+1}, \left(g_d - \frac{\sigma_d^2}{2} \right) + \sigma_d z_d^t \right) = -\kappa \text{cov}(w_t - d_t, \sigma_d z_d^t).$$

Well, $w_t = w_{t-1} + g_w - \frac{\sigma_w^2}{2} - \kappa (w_{t-1} - d_{t-1} - \overline{wd}) + \sigma_w z_w^t$, so it is orthogonal to z_d^t .

We therefore have $\kappa \text{cov}(d_t, \sigma_d z_d^t) = \kappa \text{cov}(\sigma_d z_d^t, \sigma_d z_d^t) = \kappa \text{var}(d_t)$

and so $\beta = \frac{\kappa \text{var}(d_t)}{\text{var}(d_t)} = \kappa$

Regression 3: ΔP_{t+1}^j on ΔS_{t+1} (the one-year return on a wage bond of maturity j)

In logs, $P_t^j = ((1-\kappa)^j - 1)(w_t - d_t - \overline{wd}) + \left(c_{j,w} \tilde{g}_w + c_{j,d} \tilde{g}_d + c_{j,\sigma_w} \frac{\sigma_w^2}{2} + c_{j,\sigma_d} \frac{\sigma_d^2}{2} - jr \right)$.

The terms in the second set of parentheses do not depend on w_t or d_t , so returns are

$P_{t+1}^{j-1} - P_t^j = \text{const} + ((1-\kappa)^{j-1} - 1)(w_{t+1} - d_{t+1}) - ((1-\kappa)^j - 1)(w_t - d_t)$.

$\text{cov}(P_{t+1}^{j-1} - P_t^j, \Delta d_{t+1}) = \text{cov}(((1-\kappa)^{j-1} - 1)(w_{t+1} - d_{t+1}) - ((1-\kappa)^j - 1)(w_t - d_t), \text{const} + \sigma_d z_d^{t+1})$

But w_{t+1} , w_t , d_t are all orthogonal to $\sigma_d z_d^{t+1}$, so we have $((1-\kappa)^{j-1} - 1) \text{cov}(d_{t+1}, \sigma_d z_d^{t+1})$

$= ((1-\kappa)^{j-1} - 1) \text{var}(\sigma_d z_d^{t+1})$

$\Rightarrow \beta = ((1-\kappa)^{j-1} - 1)$

C. Measuring actuarial imbalance using future taxes and benefit flows

To compute both the actuarial and market value of the unfunded obligation measures we must determine the tax and benefit cash flows associated with each cohort. From the perspective of an individual retiree, Social Security benefits are a life annuity whose magnitude depends on her history of wages (relative to the average wage in the economy), the date at which she retires, and most importantly for us, the aggregate wage when she reaches age 60. Expressing benefits as wage bonds thus requires attributing benefits to each cohort separately. We use data from the Office of the Chief Actuary (OACT) of the Social Security Administration to assign cash flows to each year and cohort as needed to determine when the uncertainty of payments is resolved.

The OACT has provided us the following 2009 Trustees Report projections for each of the next 75 years: 1) OASI payroll taxes and taxes on benefits for the Open Group, Closed Group, and Maximum Transition Cost; 2) the number of retired workers by cohort and claiming age; and 3) the average retired worker's benefit by cohort and claiming age. These projections from the 2009 Trustees Report (which the OACT derives from a microsimulation using detailed demographic, economic and program data) offer the best prediction of the future cash flows of the Social Security system. In

addition we have received the OACT estimates of the three measures of unfunded obligations related specifically to OASI (whereas only OASDI estimates appear in regularly published actuarial notes). We use this data to construct actuarial measures for retired workers receiving benefits based on their own earnings record, the population that accounts for the vast majority of Social Security benefits and for whom benefits are clearly tied to average wages at age 60. Each measure is the difference between the present value of the relevant benefit payments and tax revenues for that measure minus the \$2.4 trillion Social Security Trust Fund.

In constructing all three of our measures we take OASI payroll tax revenue flows for that measure as given. It is unnecessary to allocate these annual tax flows to particular cohorts because they are related to wages in the year they are paid. Up to an income cap, payroll tax flows in any future year will be proportional to aggregate income below that cap. Wage bond prices will capture all priced risk in these flows, so long as we assume the proportion of wages below the cap is constant (or that variation in that proportion is not associated with any market priced risk) and the size of the working population is orthogonal to market-priced risk. For the infinite horizon we project tax revenue growth at a constant real rate of 1.5%, the growth rate in years 70 to 75 of the OACT data. We ignore taxation of benefits, which depends on an individual's current income as well as benefits paid to the individual, and hence depends on wages in multiple years. These taxes are projected to grow slowly from 3.6% to 5.3% of benefits paid to the OASI Open Group.

For Open Group benefits we follow a three-step procedure. First, for each future year we multiply the OACT projections for number of retired workers in each cohort by the average benefit per worker in the cohort to obtain that year's total projected benefits by cohort. We do this separately for men and women and then combine the two. Second, we obtain present values by discounting the nominal flows using Social Security projections for the Consumer Price Index and an annual real discount rate of 2.9%. Third, we scale up these benefits from what is essentially an OAI figure to one that is more consistent with our OASI tax revenues. The difference represents benefits paid to auxiliary beneficiaries like survivors and spouses based on the earnings record of the ("primary beneficiary") worker. We do this by simply increasing all flows by the percentage needed to equate the present value of our Open Group actuarial imbalance

measure to the value the OACT gave us for OASI (discounted by half a year to match the timing implied by our discount series), and this scale-up factor turns out to be about 9.2%. This gives us benefits by cohort for the Open Group, and for the infinite horizon we project growth in total real benefits of 1.8%, the growth rate over years 70 to 75, distributed between age cohorts by the average of the distributions over these years.

The final step is to determine which benefits to attribute to the more restrictive Closed Group and Transition Cost measures. For the Closed Group we simply exclude benefits paid to cohorts with ages below 18 at the time of valuation, in keeping with the procedure followed by the OACT (Wade et. al. 2009). A more complicated adjustment is needed to calculate the Transition Costs, which require an estimate of what benefits each cohort has accrued as of the transition date. Potential accrual rules for this measure are described in detail in Geanakoplos and Zeldes (2010). We obtain an estimate of the percentage of benefits each cohort has accrued using the Social Security Public Use (SSPU) dataset, a 1% sample of current retirees' benefits and income histories. For each observation in the data we calculate a benefit based on the worker's average earnings to-date in a particular year and then multiply the resulting benefit by the fraction $(\text{current age} - 18) / (62 - 18)$. We then compute an accrual percentage for each age that represents the average ratio of benefits accrued by that age to the total amount accrued by cohorts that have reached age 62. Implicitly, we are assuming the profile of accruals is similar across cohorts. For the Immediate Transition Cost we multiply each Closed Group cohort's benefits by the relevant percentage for its age in a particular year to obtain the amount of its benefits the cohort has earned as of the immediate valuation date of December 31, 2008. For the Future Transition Cost we start from the infinite-horizon Open Group benefits and then apply our accrual percentages to the benefits received more than 75 years in the future by the cohorts that have reached age 18 by the future valuation date of December 31, 2083.

In the end our actuarial estimates do not quite match the OACT estimates for the three published measures of unfunded obligations related to OASI. Our Open Group measure matches by construction, but we obtain unfunded obligations of \$13.7 trillion (compared to \$14.7) for the Closed Group and \$14.2 trillion (vs. \$16.5) for the Maximum Transition Cost. In part this reflects our use of a constant real discount rate equal to the Trustees Report intermediate cost assumption about the interest rate in all years

from 2018 onward, rather than the (unpublished) projected returns on the Trust Fund that OACT uses. There also may be variation between years or cohorts in the percentage of OASI benefits paid to workers based on their own earnings histories, i.e. our proportional increase in OAI benefits may not be appropriate for all three measures. Presumably with more information we could match the actuarial measures exactly, but this is not our intent. Our goal is to show how the cointegration risk adjustment affects these measures, and this effect appears consistent across alternative procedures we have used (Geanakoplos and Zeldes 2010).

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.
- Beaudry, Paul and Franck Portier, 2004, "An Exploration into Pigou's Theory of Cycles" *Journal of Monetary Economics* 51(6),1183-1216.
- Blocker, Alexander W., Laurence J. Kotlikoff, and Stephen A. Ross, 2008, "The True Cost of Social Security," NBER Working Paper 14427, 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
- Cox, John C., Ross, Stephen A. Ross and Mark Rubenstein 1979 "Option Pricing: A Simplified Approach" *Journal of Financial Economics* Vol 7, September: 229-263
- Davis, Steven J. and Paul Willen, 2000, "Using Financial Assets to Hedge Labor Income Risks: Estimating the Benefits," presented at NBER Program on Economic Fluctuations and Growth, February 5.
- 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, 2009, "Reforming Social Security with Progressive Personal Accounts," in *Social Security Policy in a Changing Environment*. Jeffrey R. Brown, Jeffrey B. Liebman, and David A. Wise, eds. Chicago: University of Chicago Press, 73-121.
- Geanakoplos, John, and Stephen P. Zeldes. 2010. "Market Valuation of Accrued Social Security Benefits," in *Measuring and Managing Federal Financial Risk*, Deborah Lucas, ed. Chicago: University of Chicago Press, 213-233.
- Goetzmann, William, 2008, "More Social Security, Not Less," *Journal of Portfolio Management*, vol 35, no. 1, pp. 115-123.
- 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.

- Jaimovich, Nir, and Sergio Rebelo, 2009, "Can News about the Future Drive the Business Cycle?" *American Economic Review*, 99, 1097-1118.
- Jermann, Urban J., 1999, "Social Security and Institutions for Intergenerational, Intragenerational, and International Risk-Sharing: A Comment," *Carnegie-Rochester Conference Series on Public Policy*, Vol 50, pp. 205-212.
- Khorasanee, Zaki, 2009, "What Discount Rate Should Be Used to Value a Cash-Flow Linked to Final Salary?" *Journal of Pension Economics and Finance*, Vol 8, No 3, 351-360.
- Koehler, Paul A. and Laurence J. Kotlikoff, 2009, "Marking Social Security's Open Group Liability to Market," Michigan Retirement Research Center Research Paper No. 2009-217.
- Krueger, Dirk and Felix Kubler, 2002, "Intergenerational Risk-Sharing via Social Security When Financial Markets Are Incomplete," *American Economic Review*, Vol 92, No 2, May.
- Krueger, Dirk and Felix Kubler, 2006, "Pareto-Improving Social Security Reform when Financial Markets Are Incomplete?" *American Economic Review*, Vol 96, No 3, June.
- Lucas, Deborah (ed.), 2010, *Measuring and Managing Federal Financial Risk*, Chicago: University of Chicago Press.
- 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.
- Lucas, Deborah and Stephen P. Zeldes, 2009, "How Should Public Pension Plans Invest?" *American Economic Review Papers and Proceedings*, May, 527-532.
- Novy-Marx, Robert, and Joshua Rauh, 2010, "Public Pension Liabilities: How Big Are They and What Are They Worth?", unpublished manuscript, Kellogg School, Northwestern University, forthcoming, *Journal of Finance*.
- Office of Management and Budget, 2009. "Updated Summary Tables," Budget of the U.S. Government, Fiscal Year 2010, Tables S-4, S-5.
<http://www.gpoaccess.gov/usbudget/fy10/pdf/summary.pdf>
- Santos, Tano and Pietro Veronesi, 2006, "Labor Income and Predictable Stock Returns," *The Review of Financial Studies*, Vol 19, No 1.
- Schmitt-Grohé, Stephanie and Martín Uribe, 2010, "What's News in Business Cycles," unpublished manuscript, Columbia University, April.

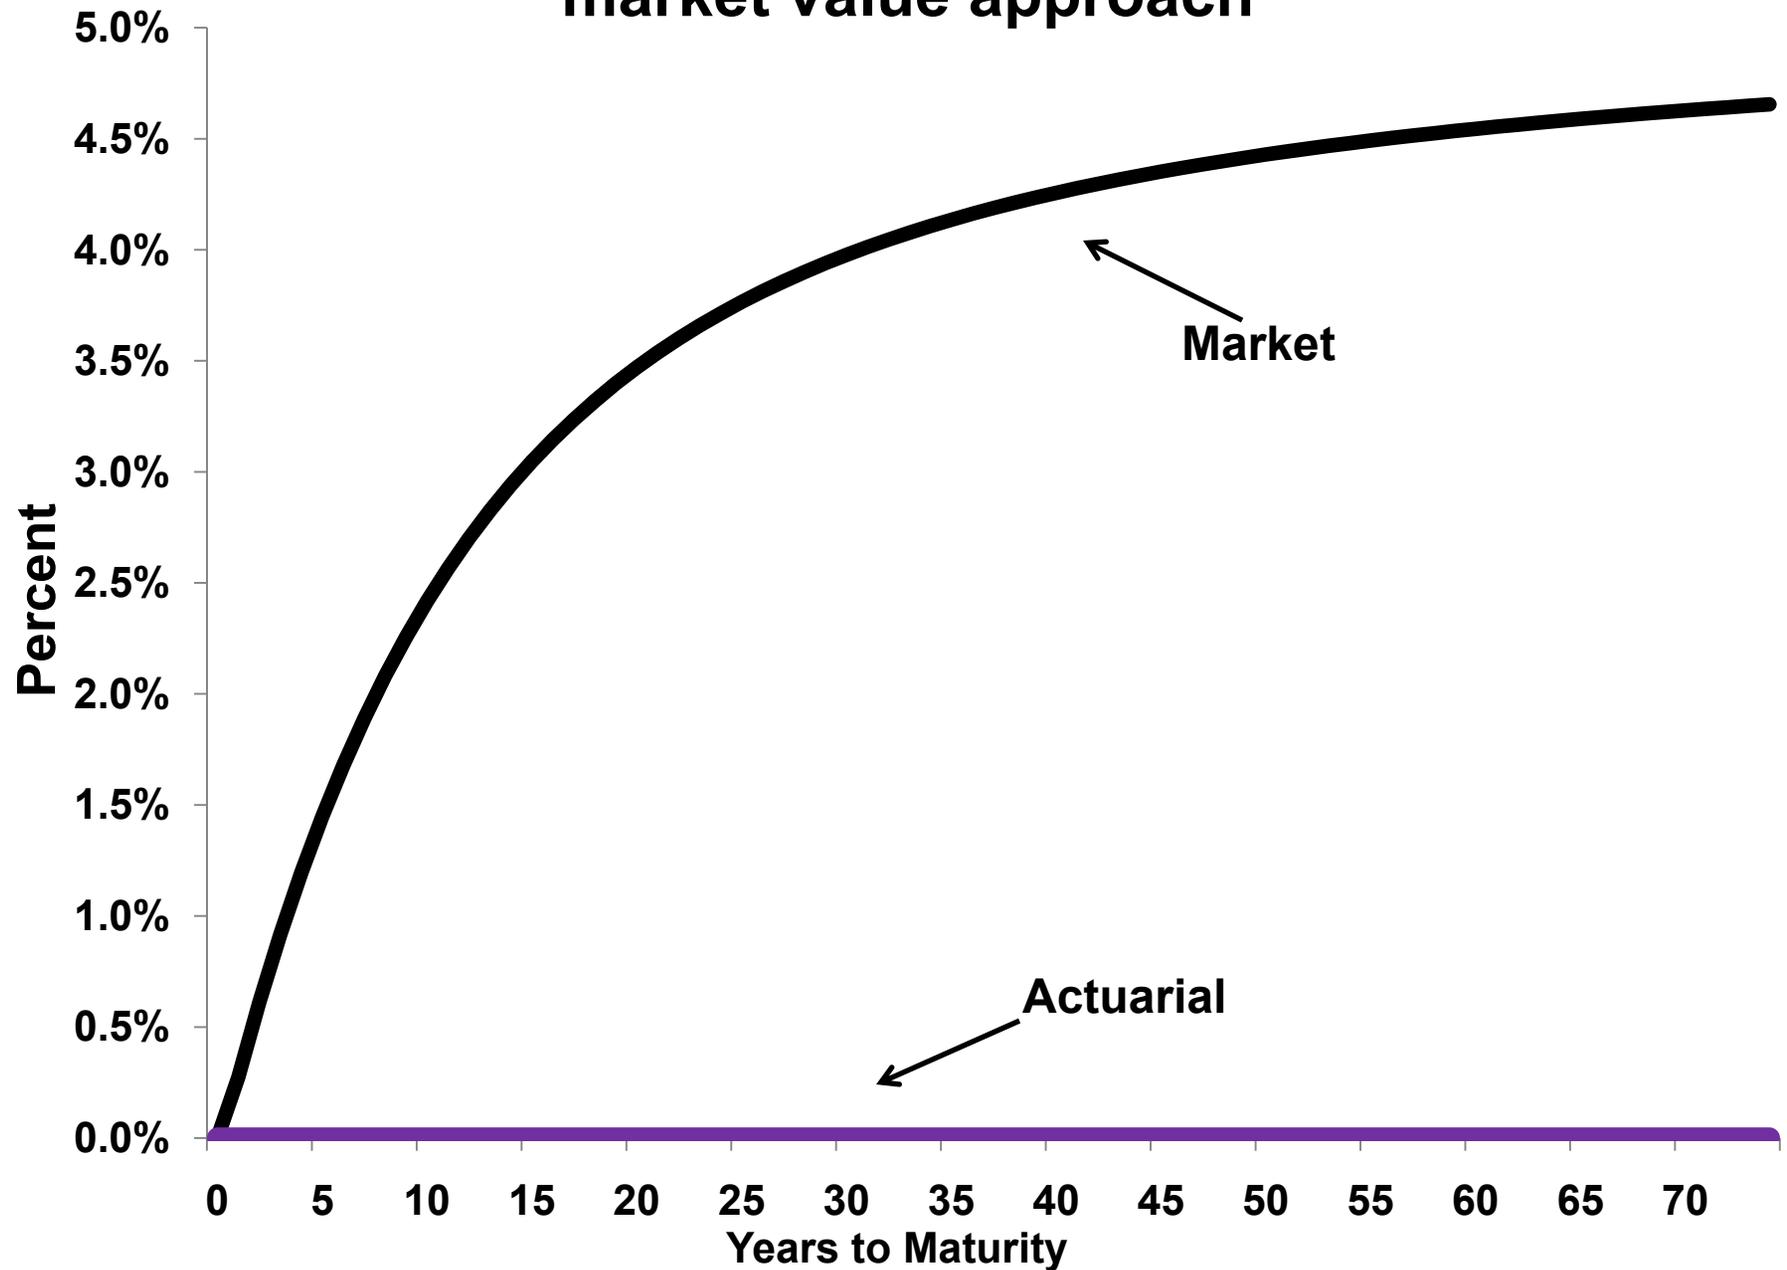
Social Security Administration, 2002. Life Tables for the United States Social Security Area 1900-2100, Actuarial Study 116, SSA Pub. No. 11-11536. Washington, DC: Office of the Chief Actuary.

Social Security Advisory Board, 2007. The Report of the 2007 Technical Panel on Assumptions and Methods. Washington, DC.

Social Security Administration, 2008. Annual Statistical Supplement to the Social Security Bulletin, 2008. Washington, DC: Office of Policy.

Wade, Alice, Jason Schultz, and Steve Goss, 2009. "Unfunded Obligation and Transition Cost for the OASDI Program," Actuarial Note #2009.1, Social Security Administration, Office of the Chief Actuary, August.

**Figure 1: Discount rate for wage bonds – risk-free rate:
market value approach**



**Figure 1b: Discount rate for wage bonds – risk-free rate:
market value approach**

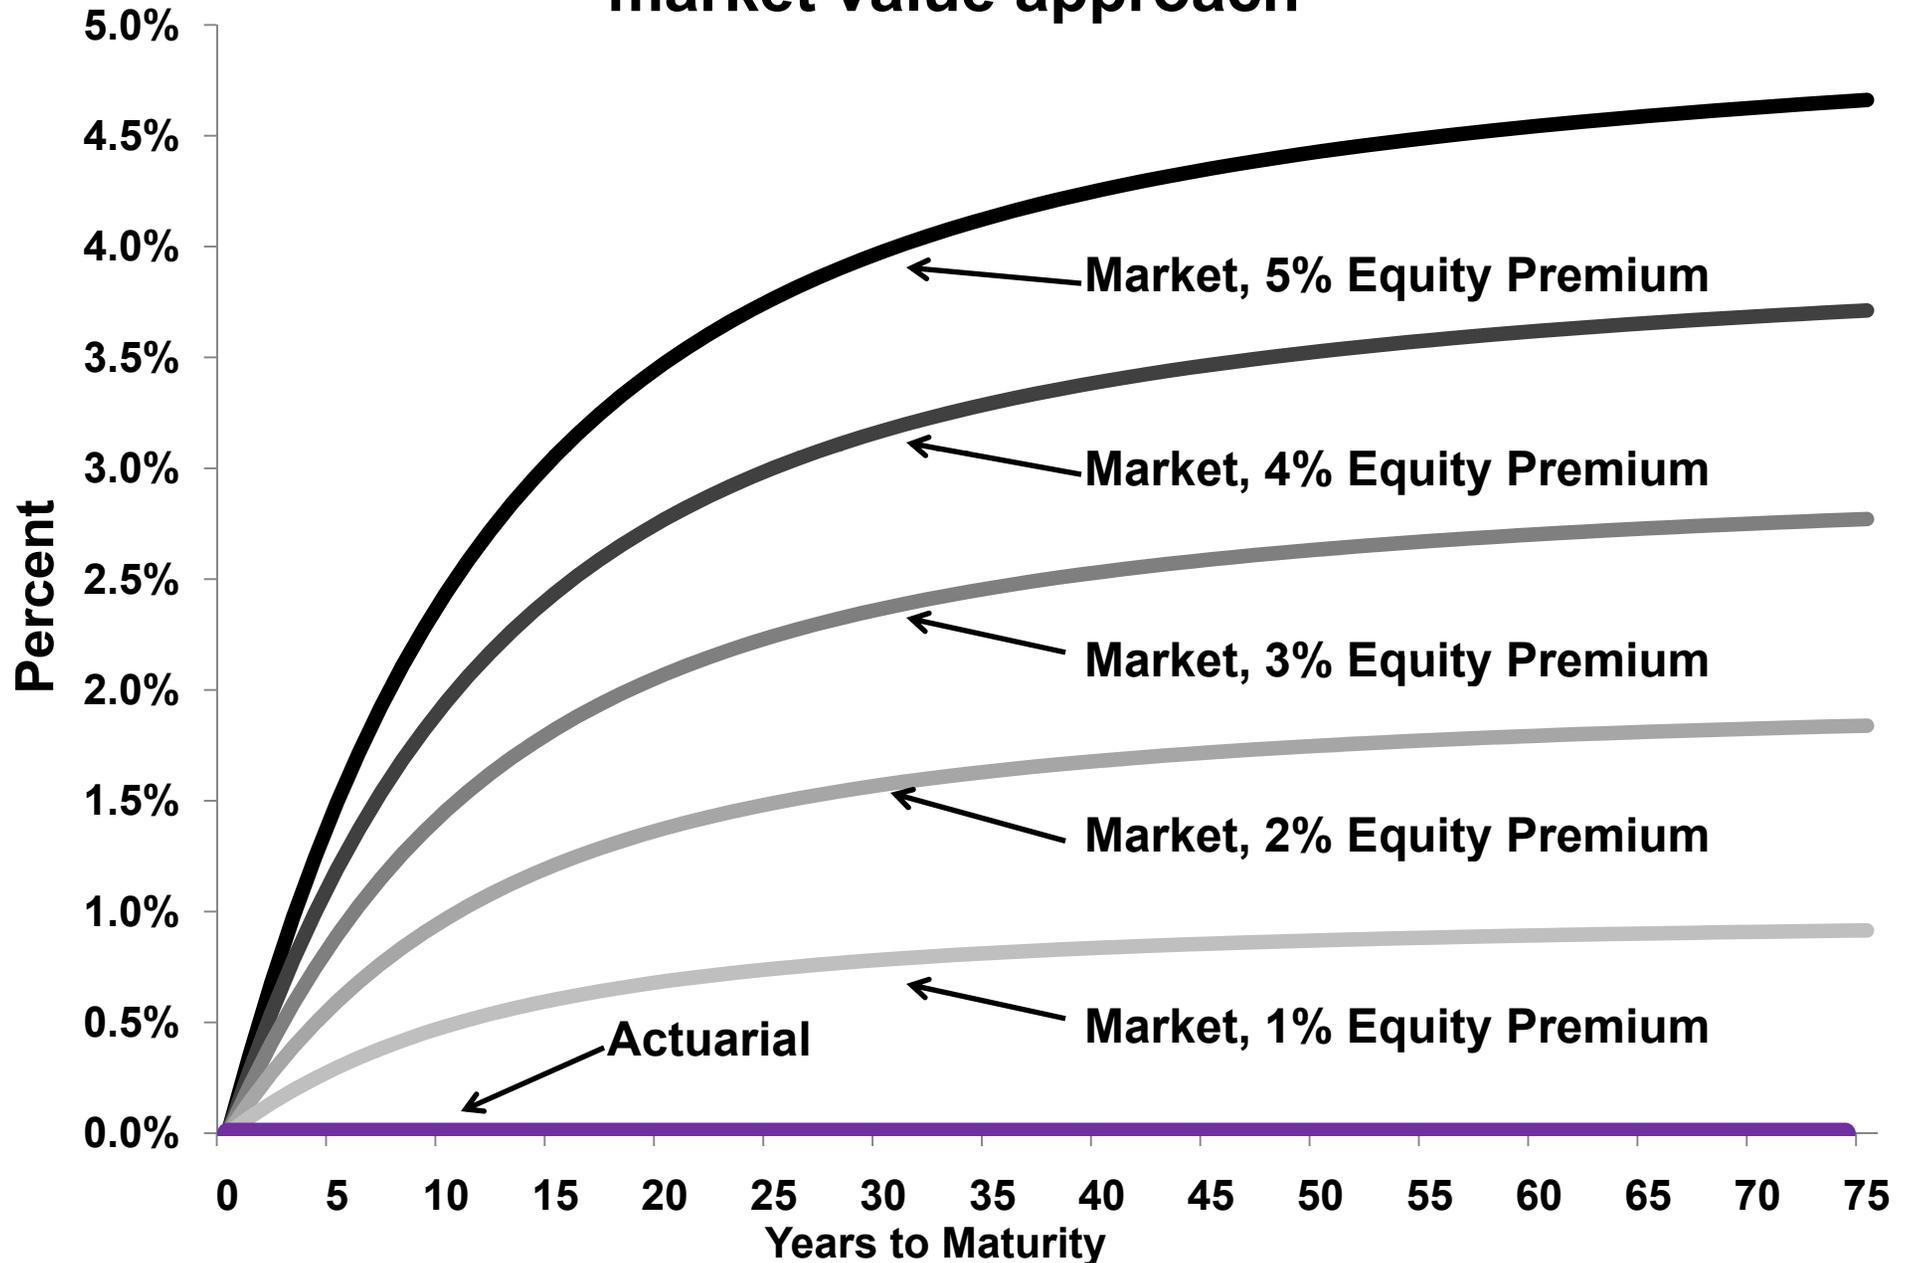
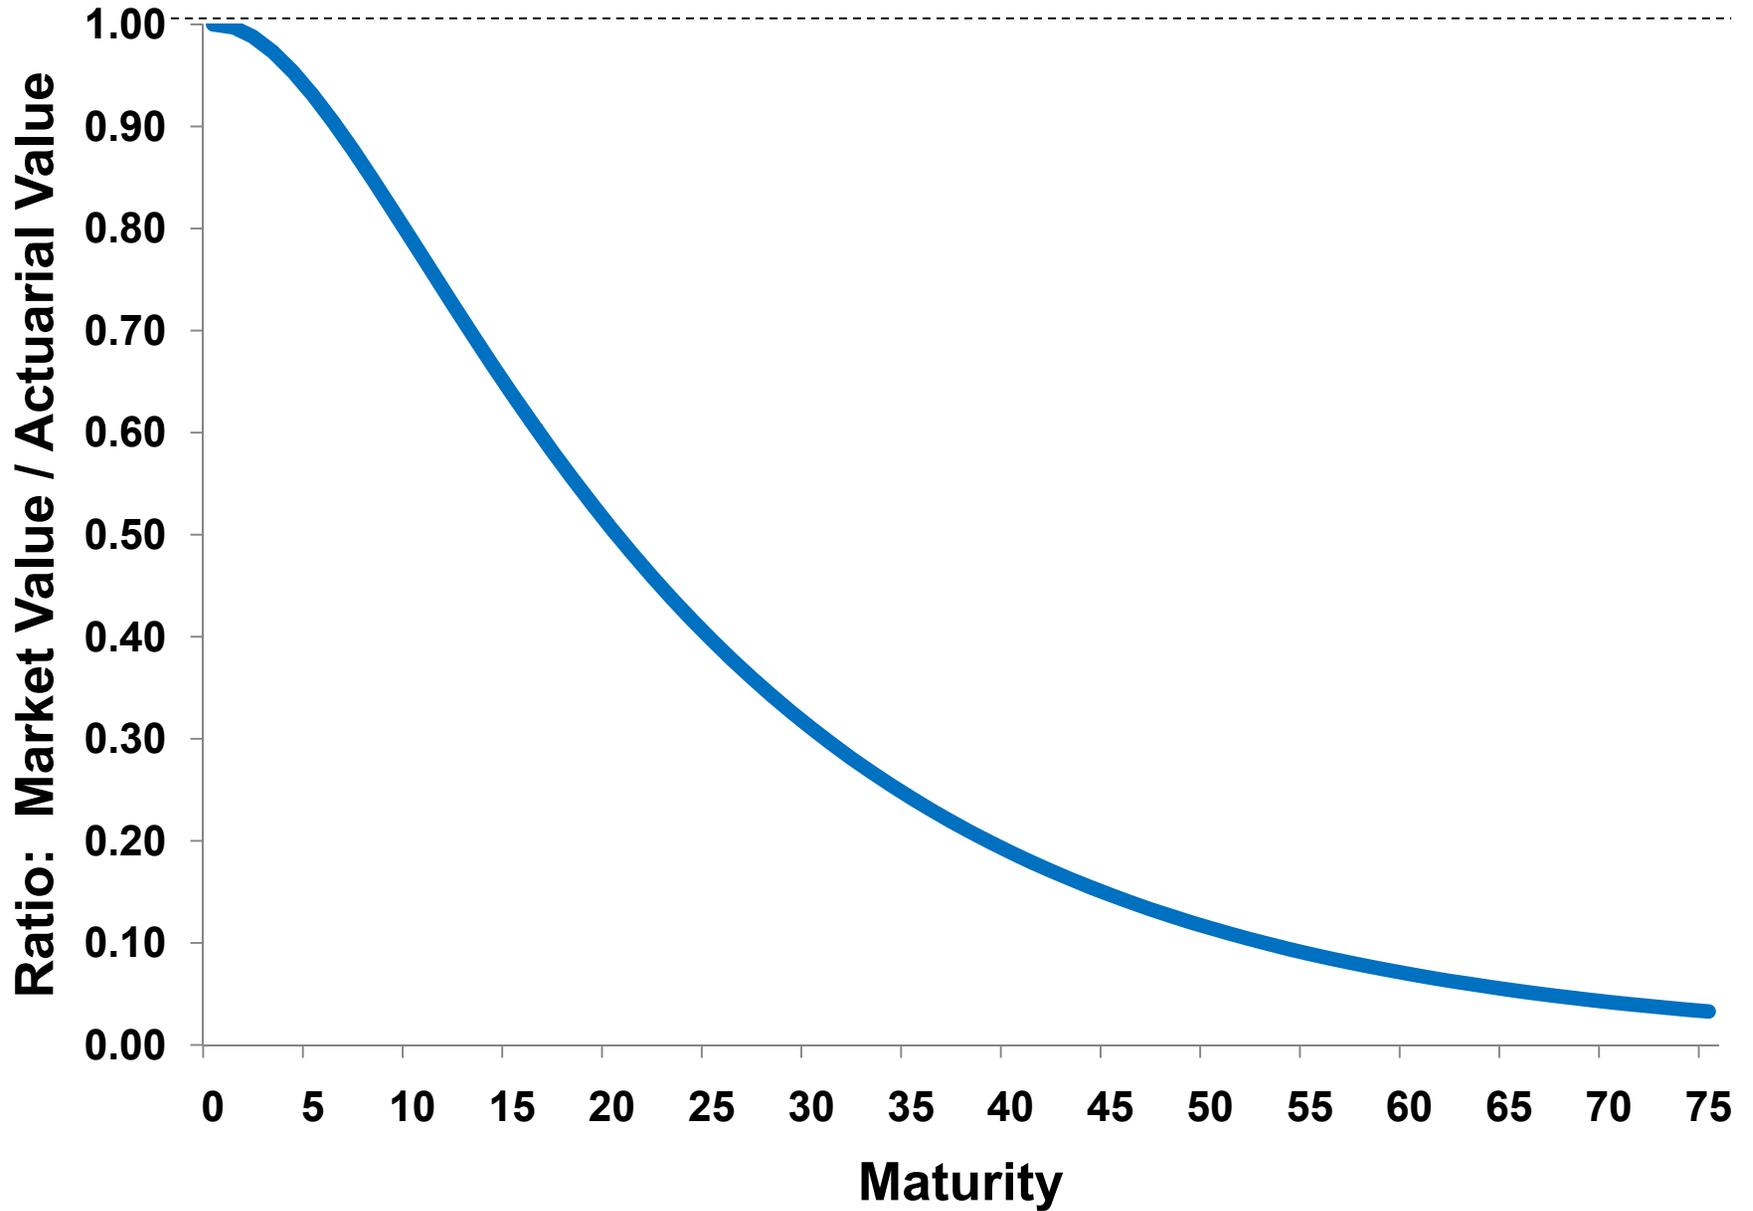


Figure 2: Ratio of Market to Actuarial Wage Bond Prices



**Figure 3: Discount rate - risk-free rate
(three approaches)**

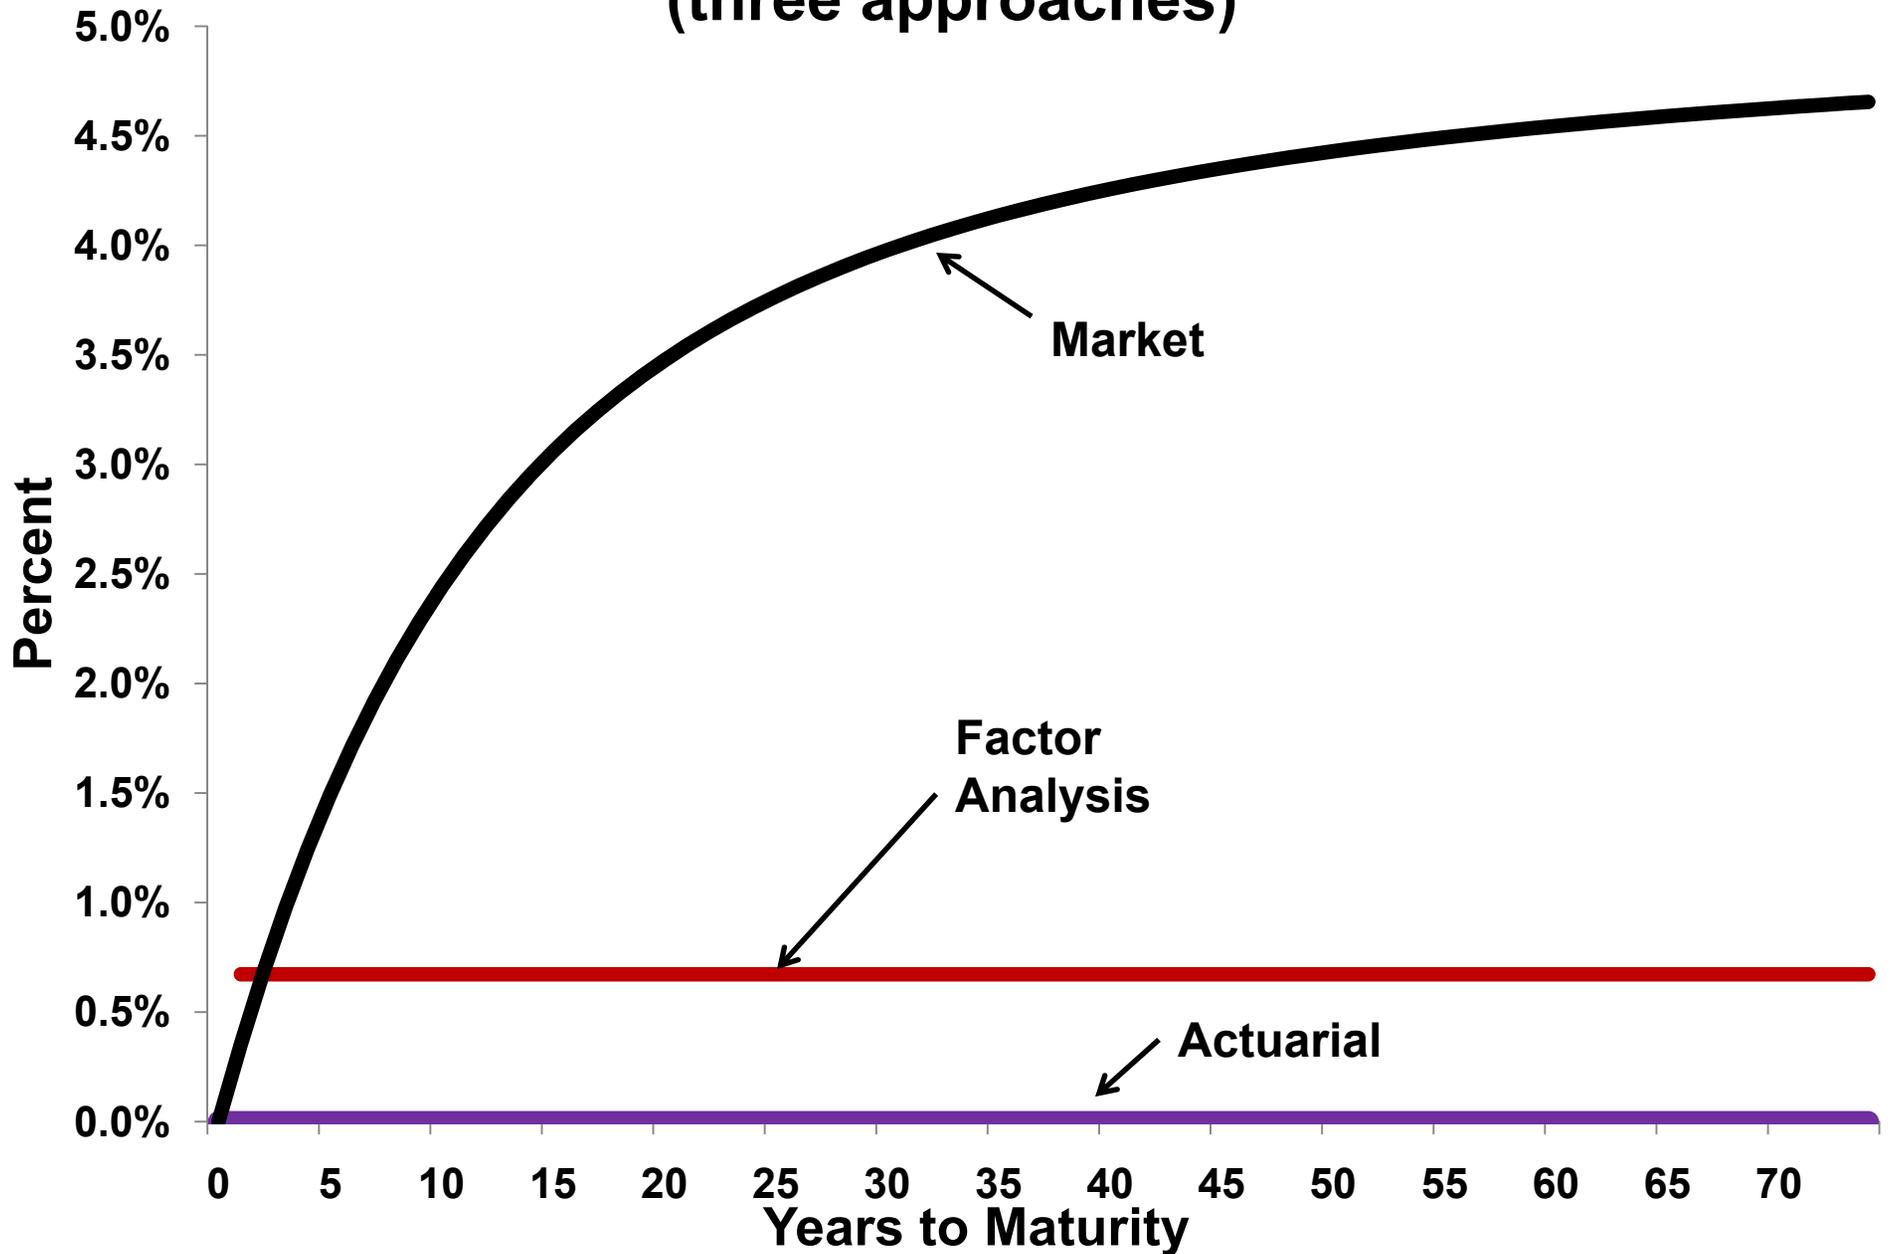


Figure 4: Real Cash Flows

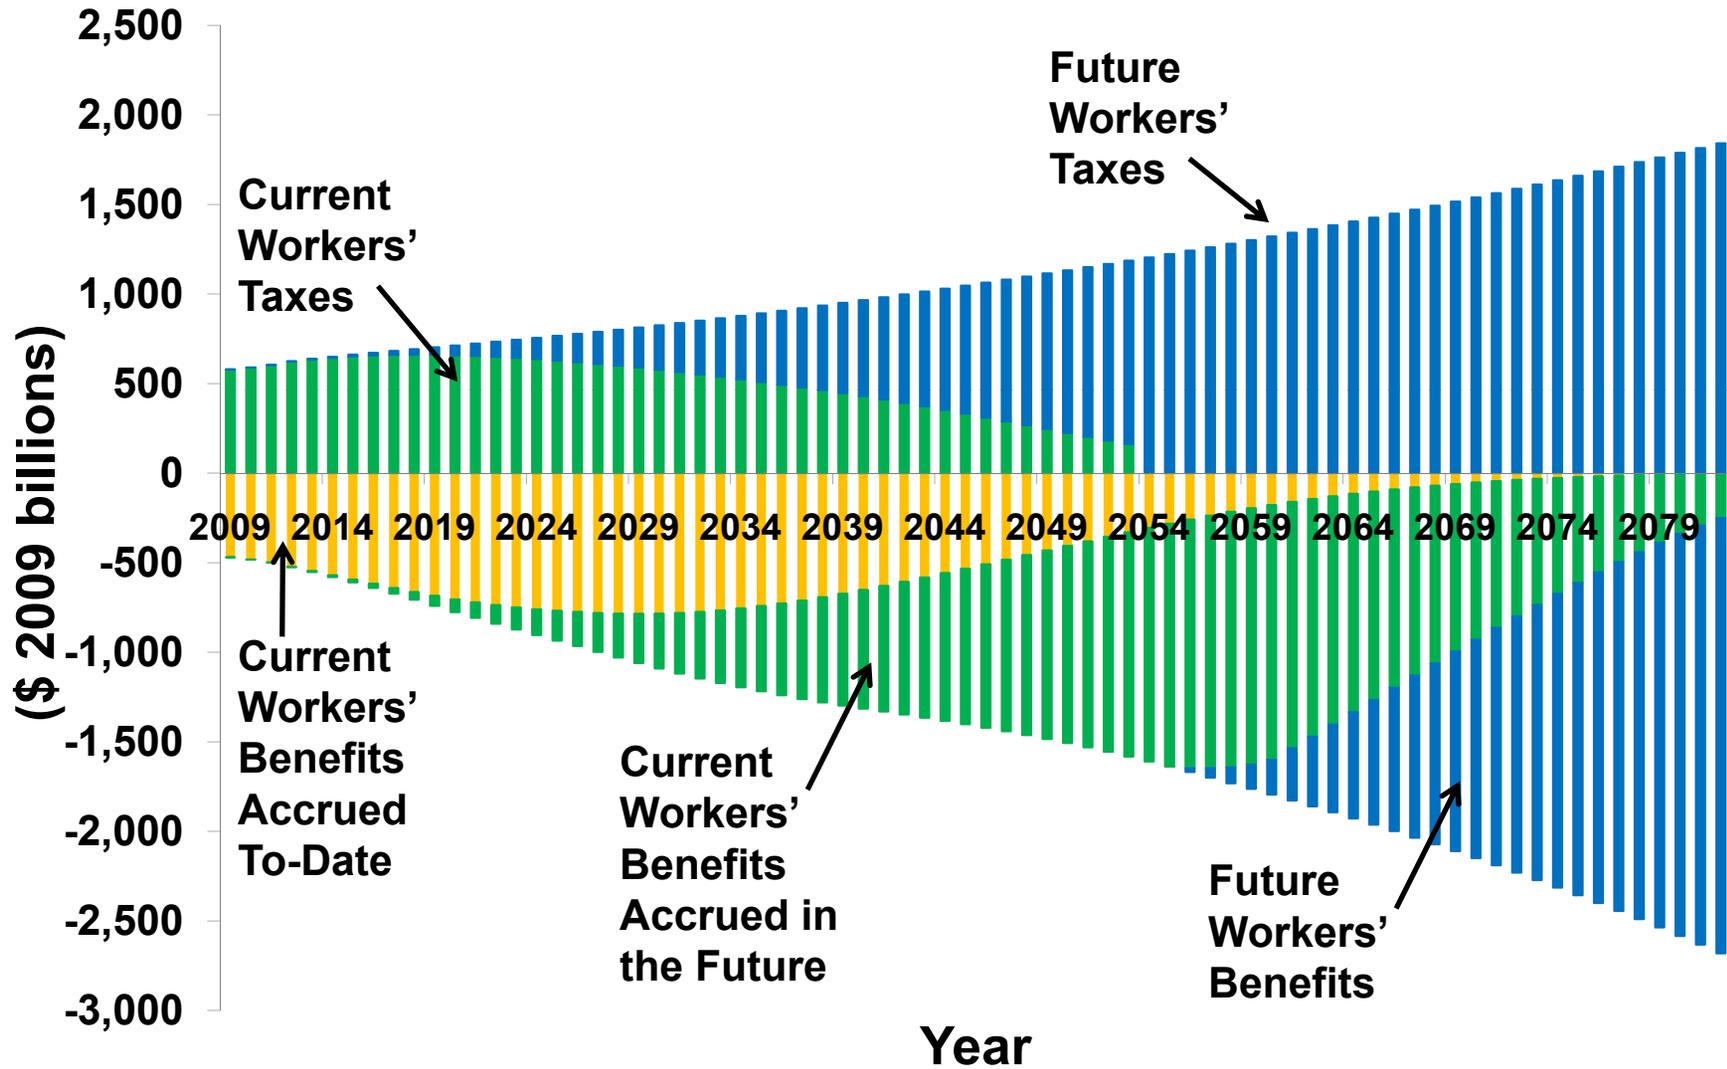


Figure 5: Annual Cash Flows, Actuarial PV

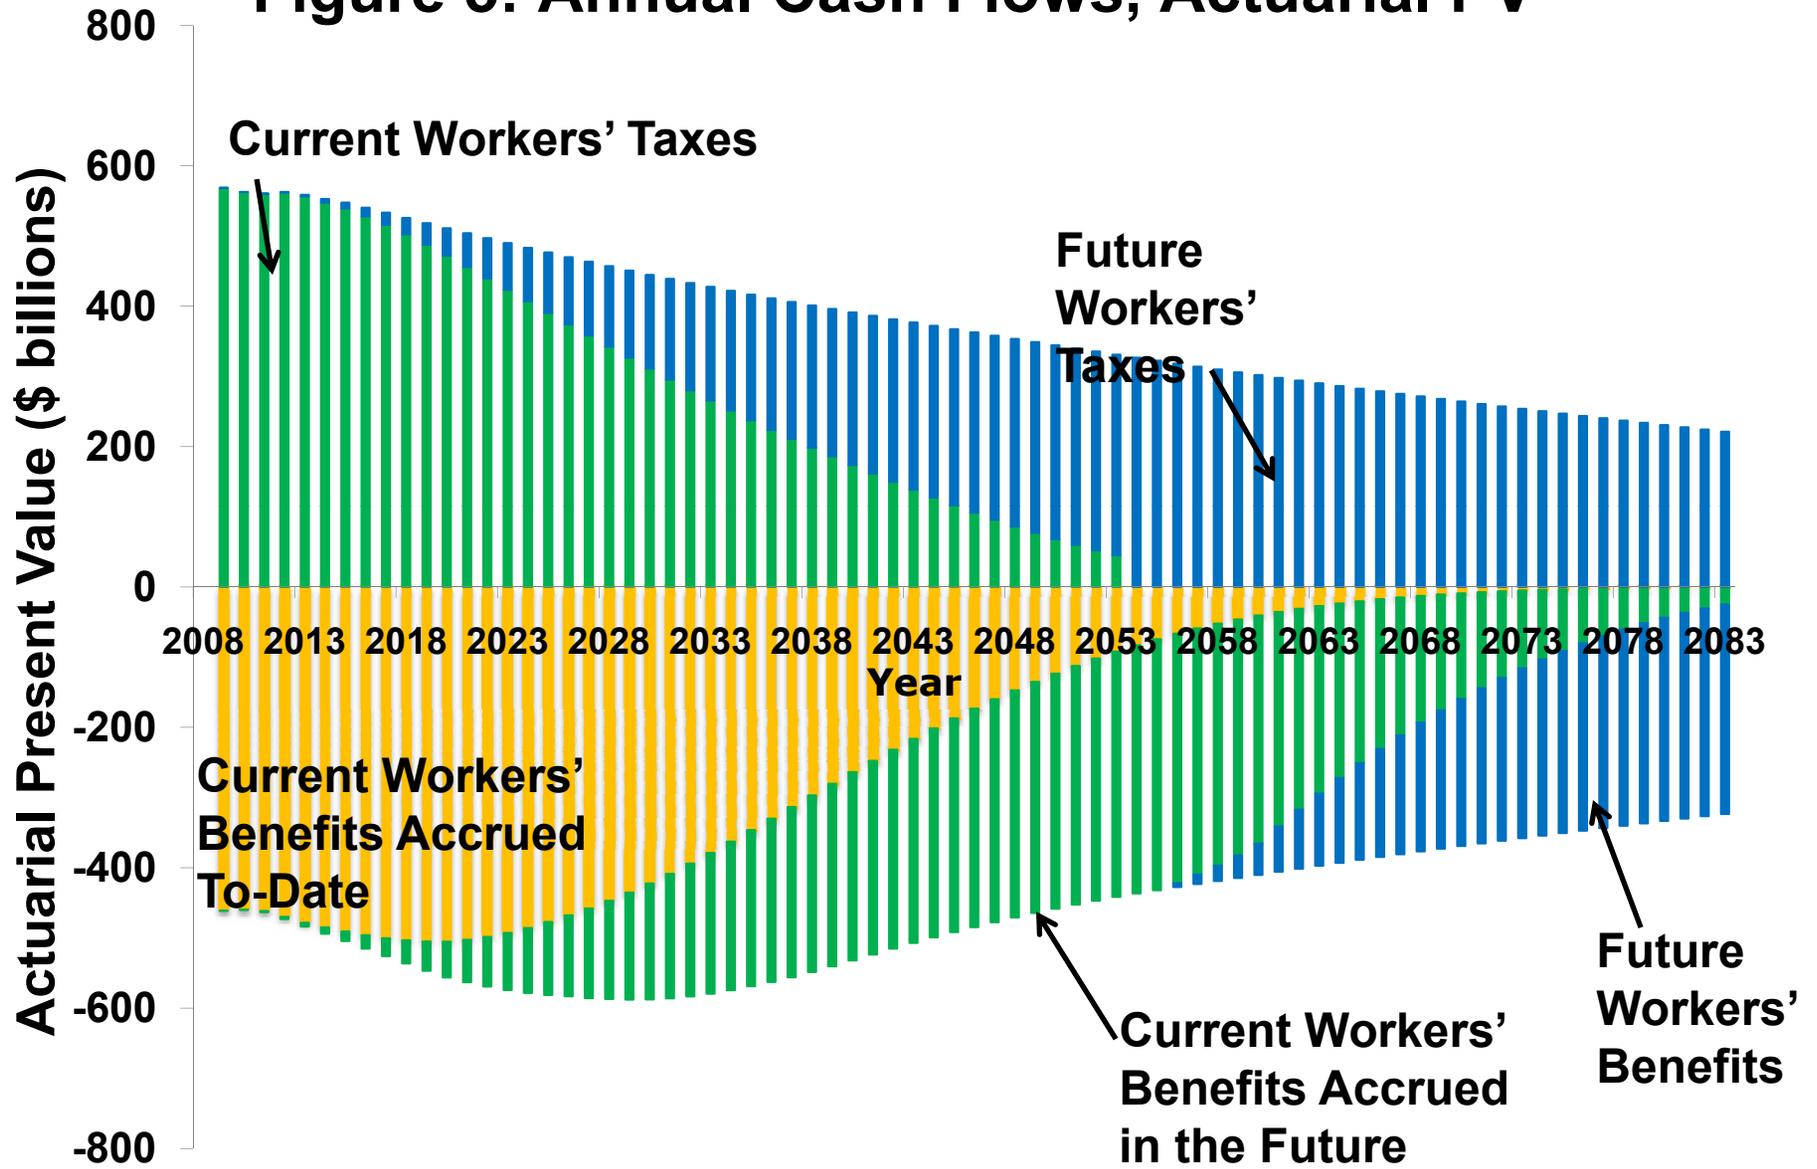


Figure 6: Real Cash Flows, Infinite Horizon

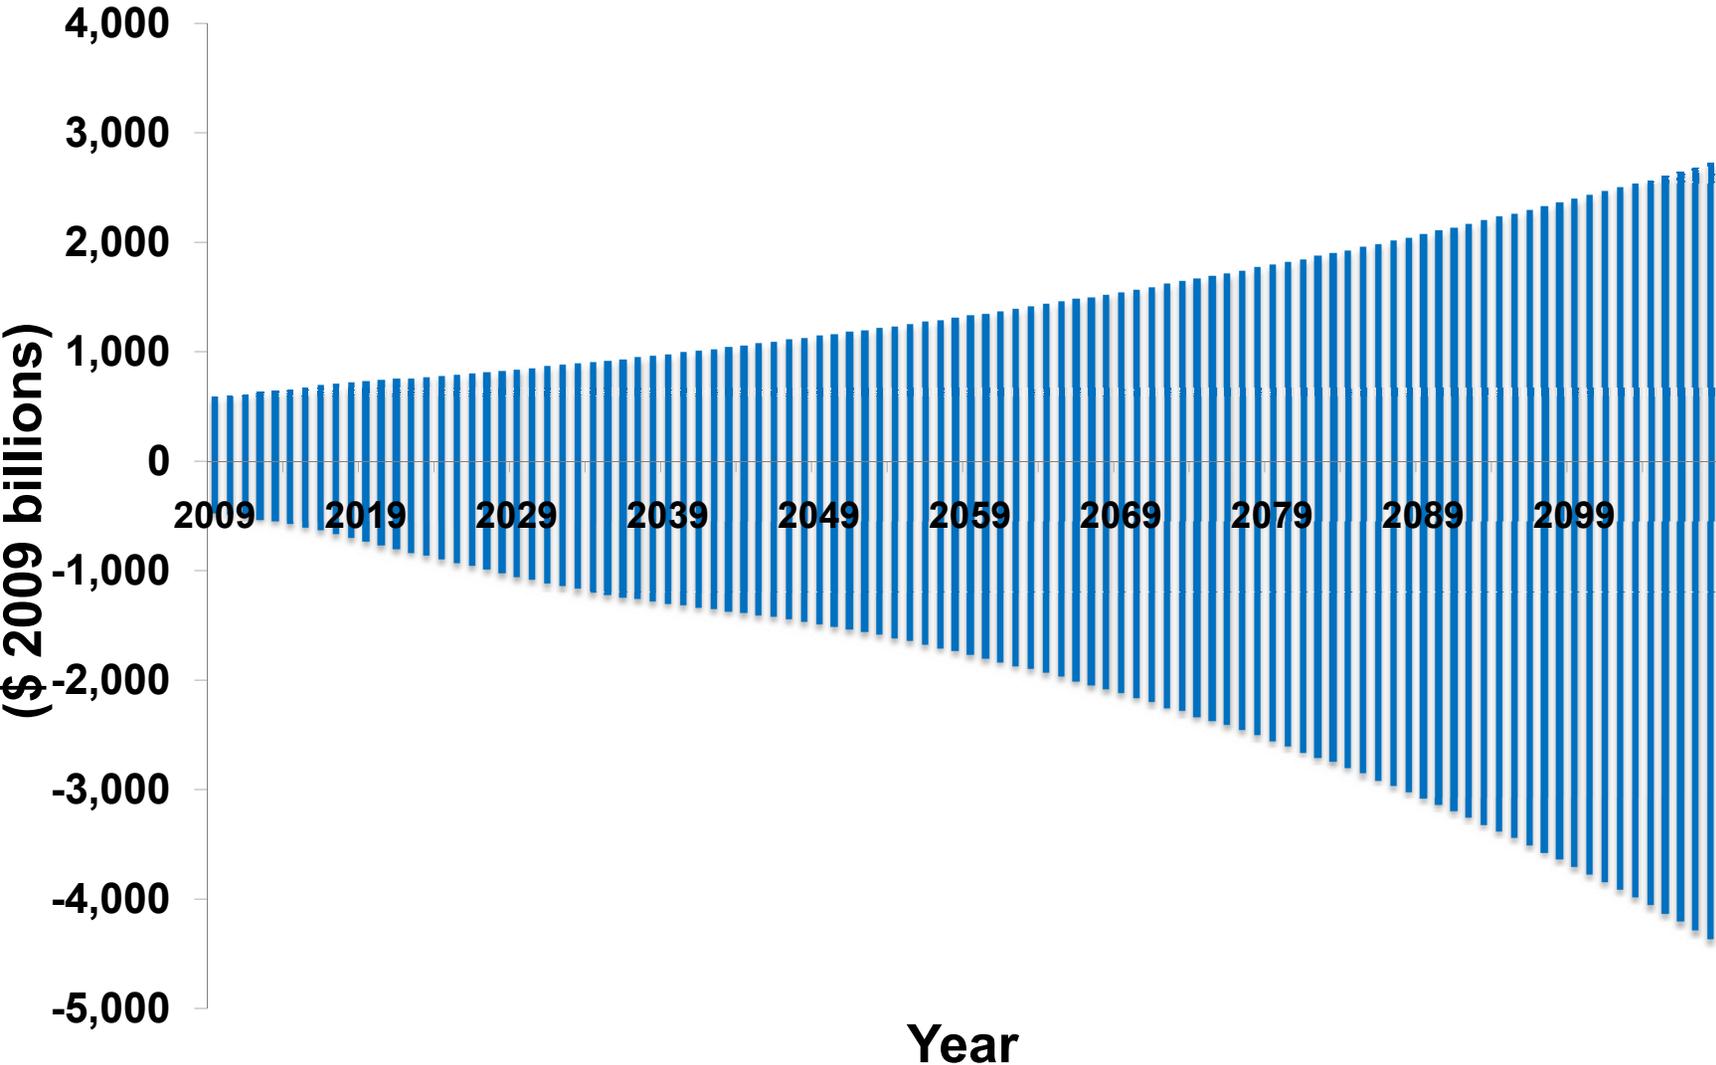


Figure 7: Value of Wage Bonds, Actuarial PV

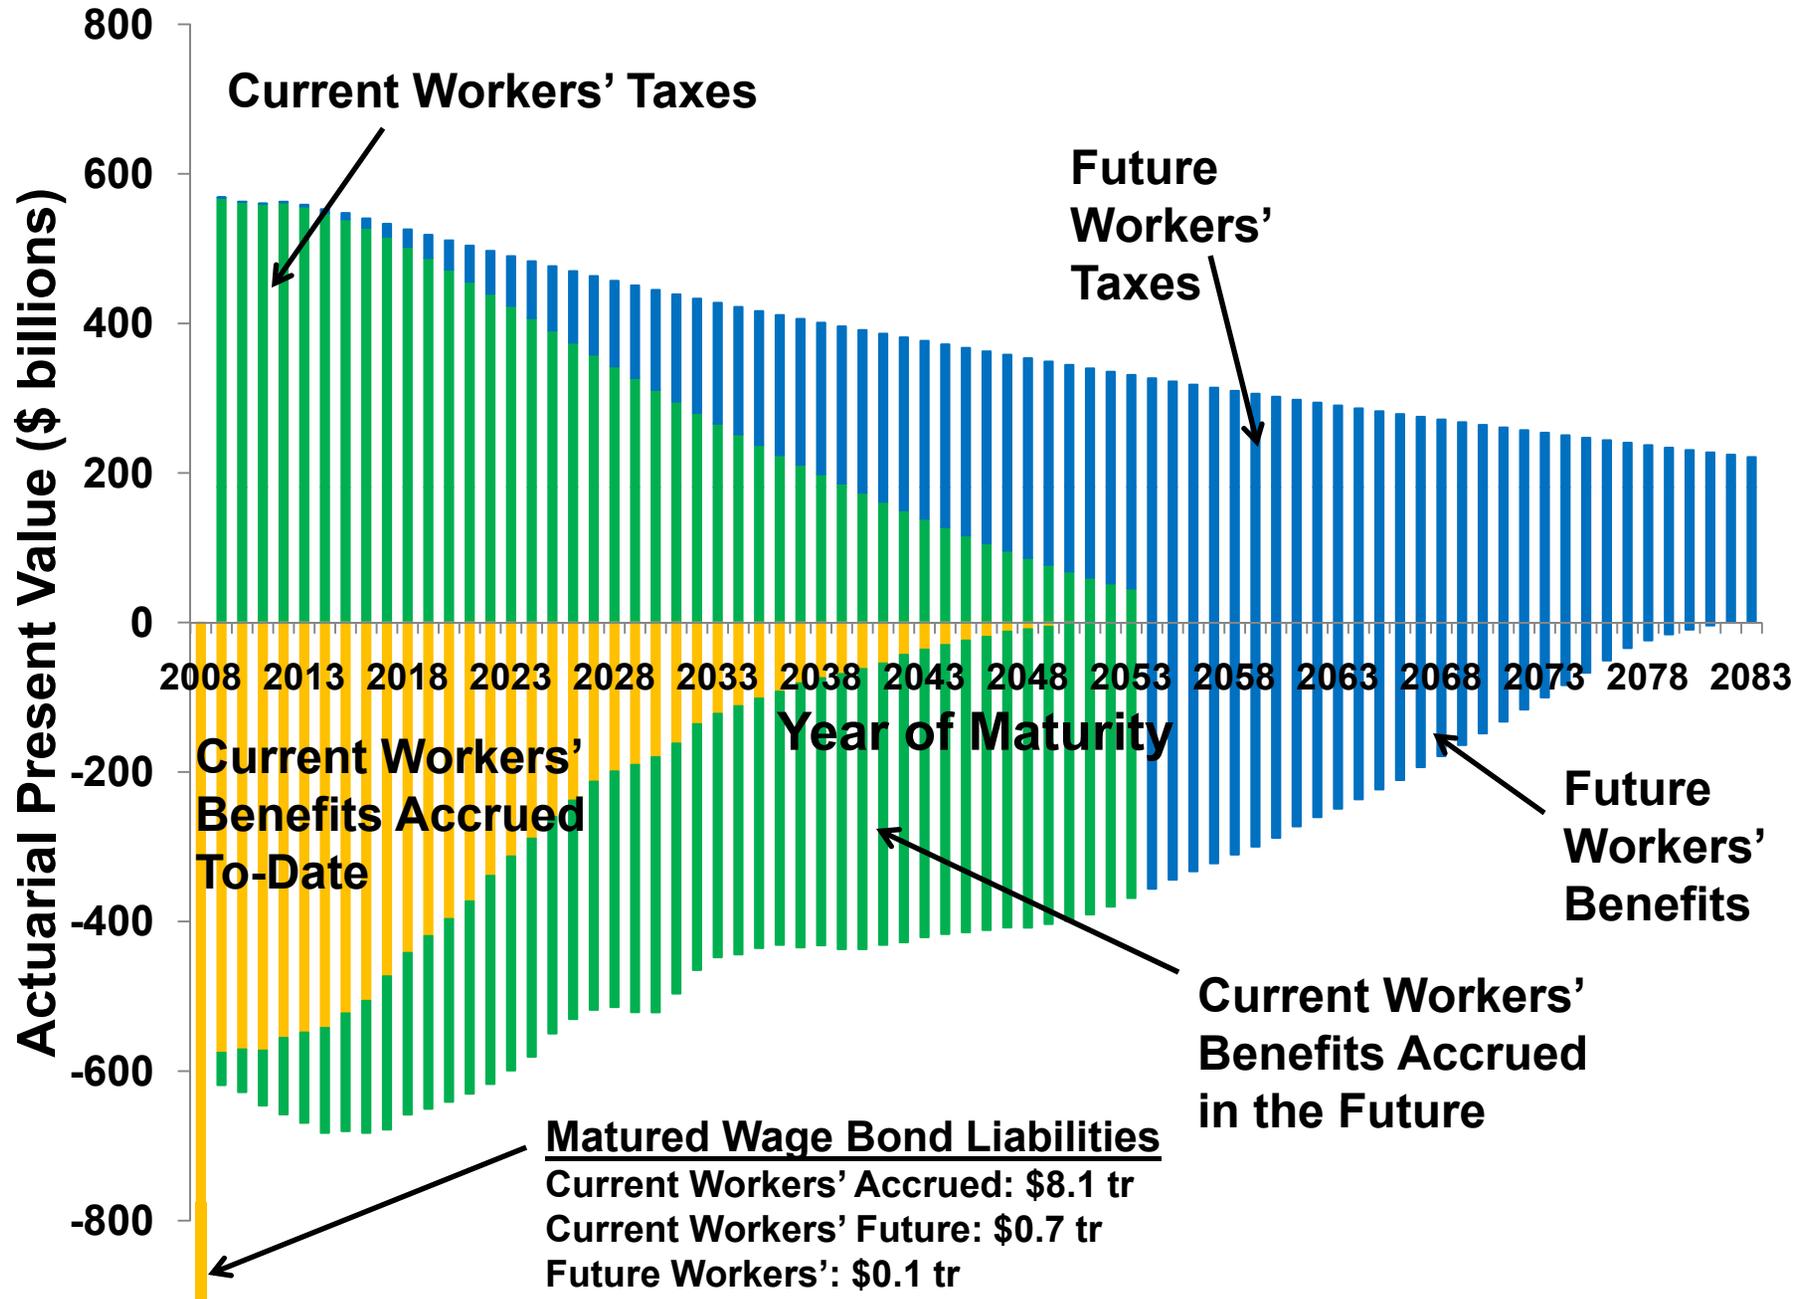


Figure 8: Net Value of Wage Bonds, Actuarial PV

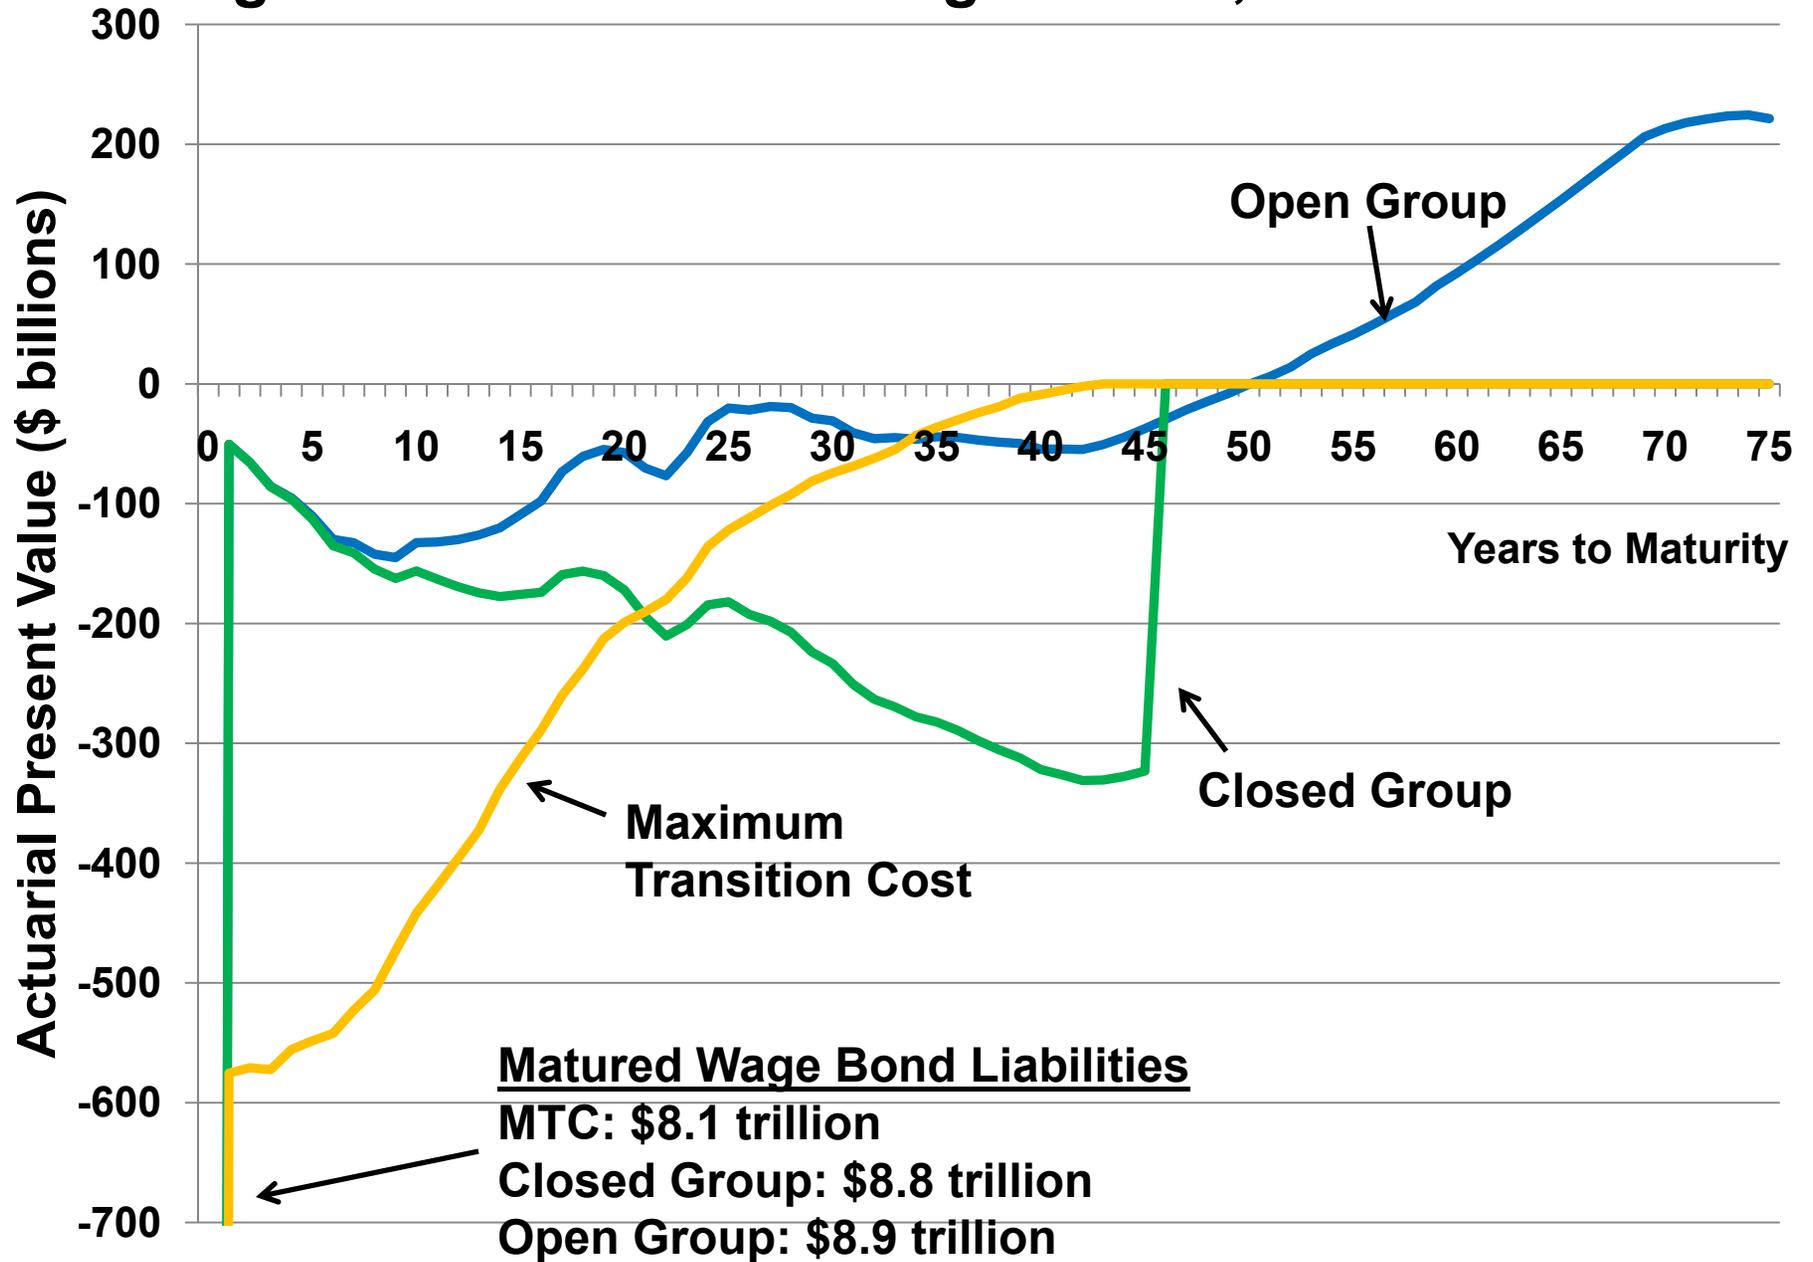


Figure 9a: Net Value of Wage Bonds, Market vs. Actuarial PV, Maximum Transition Cost

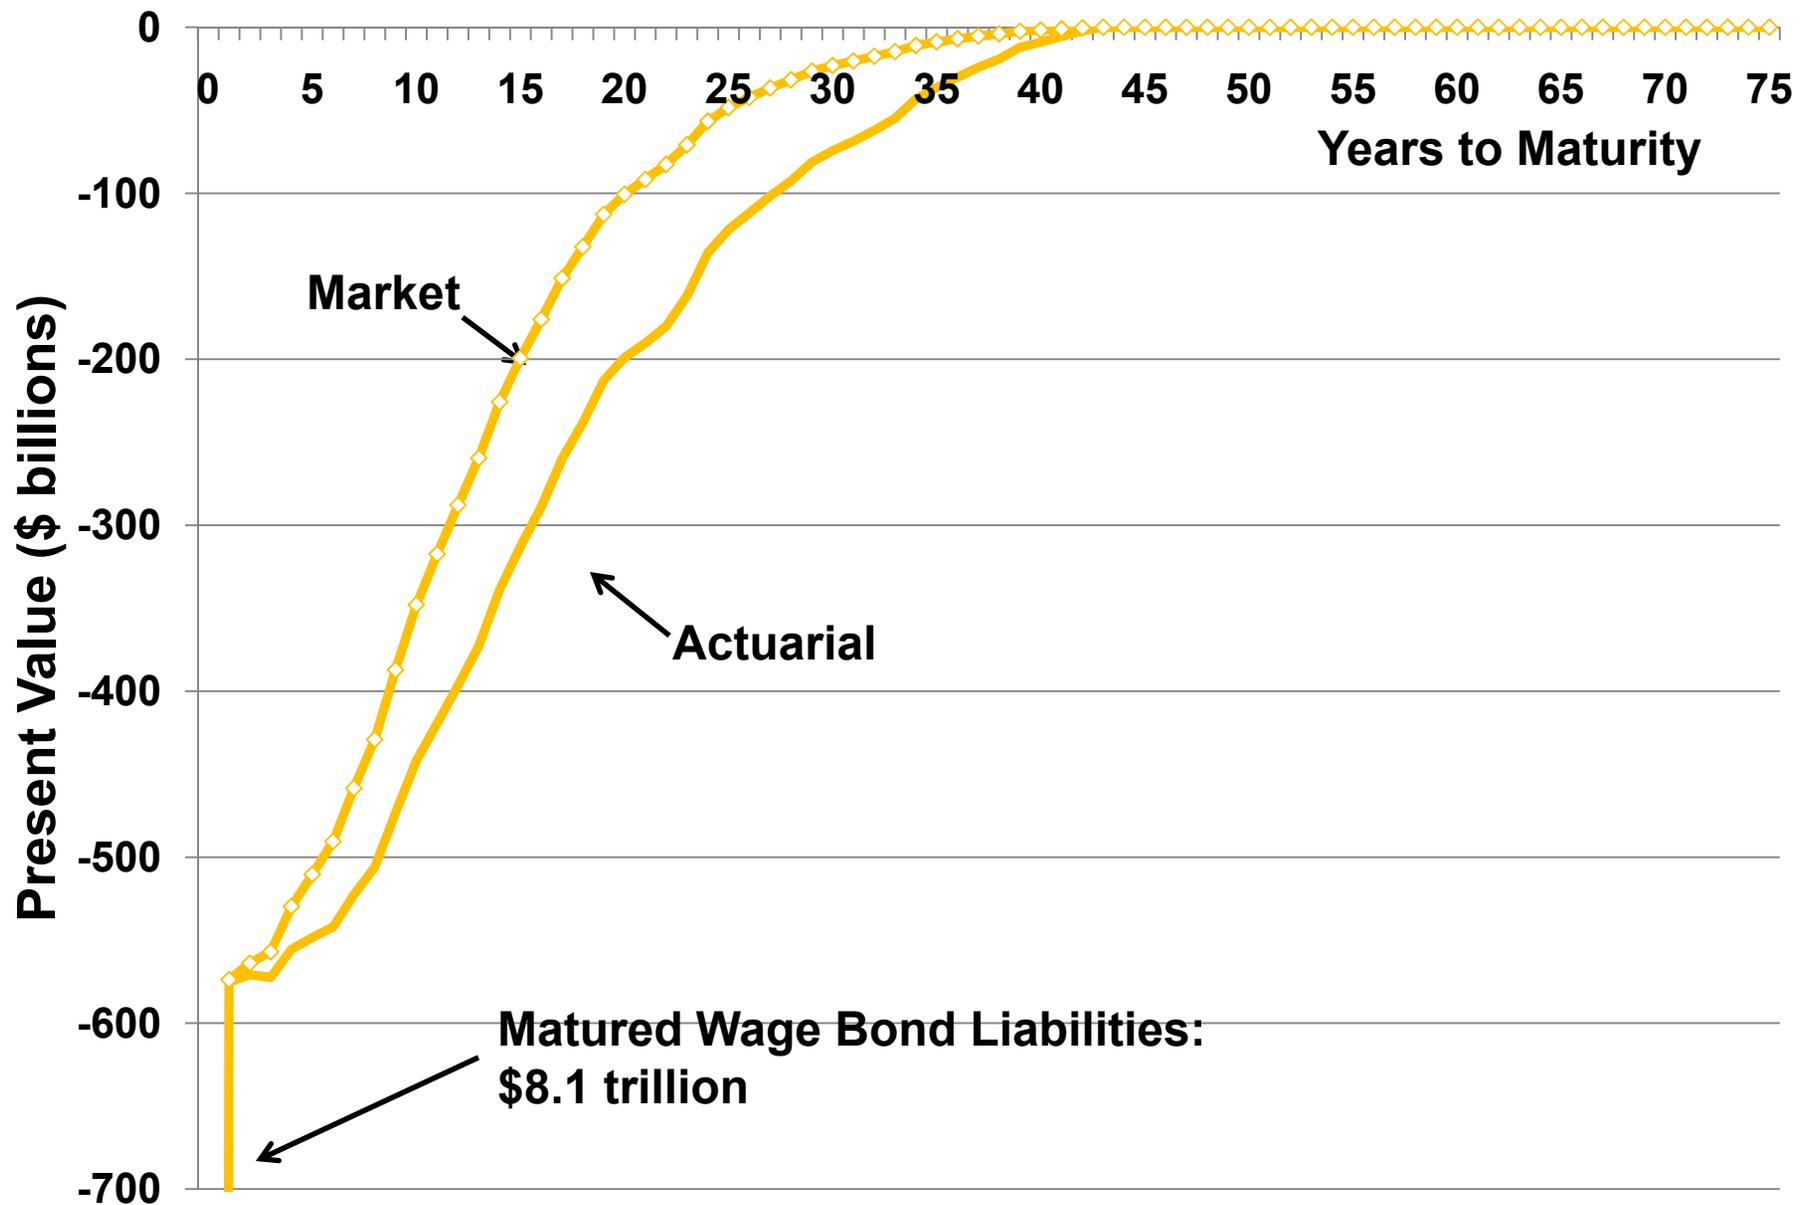
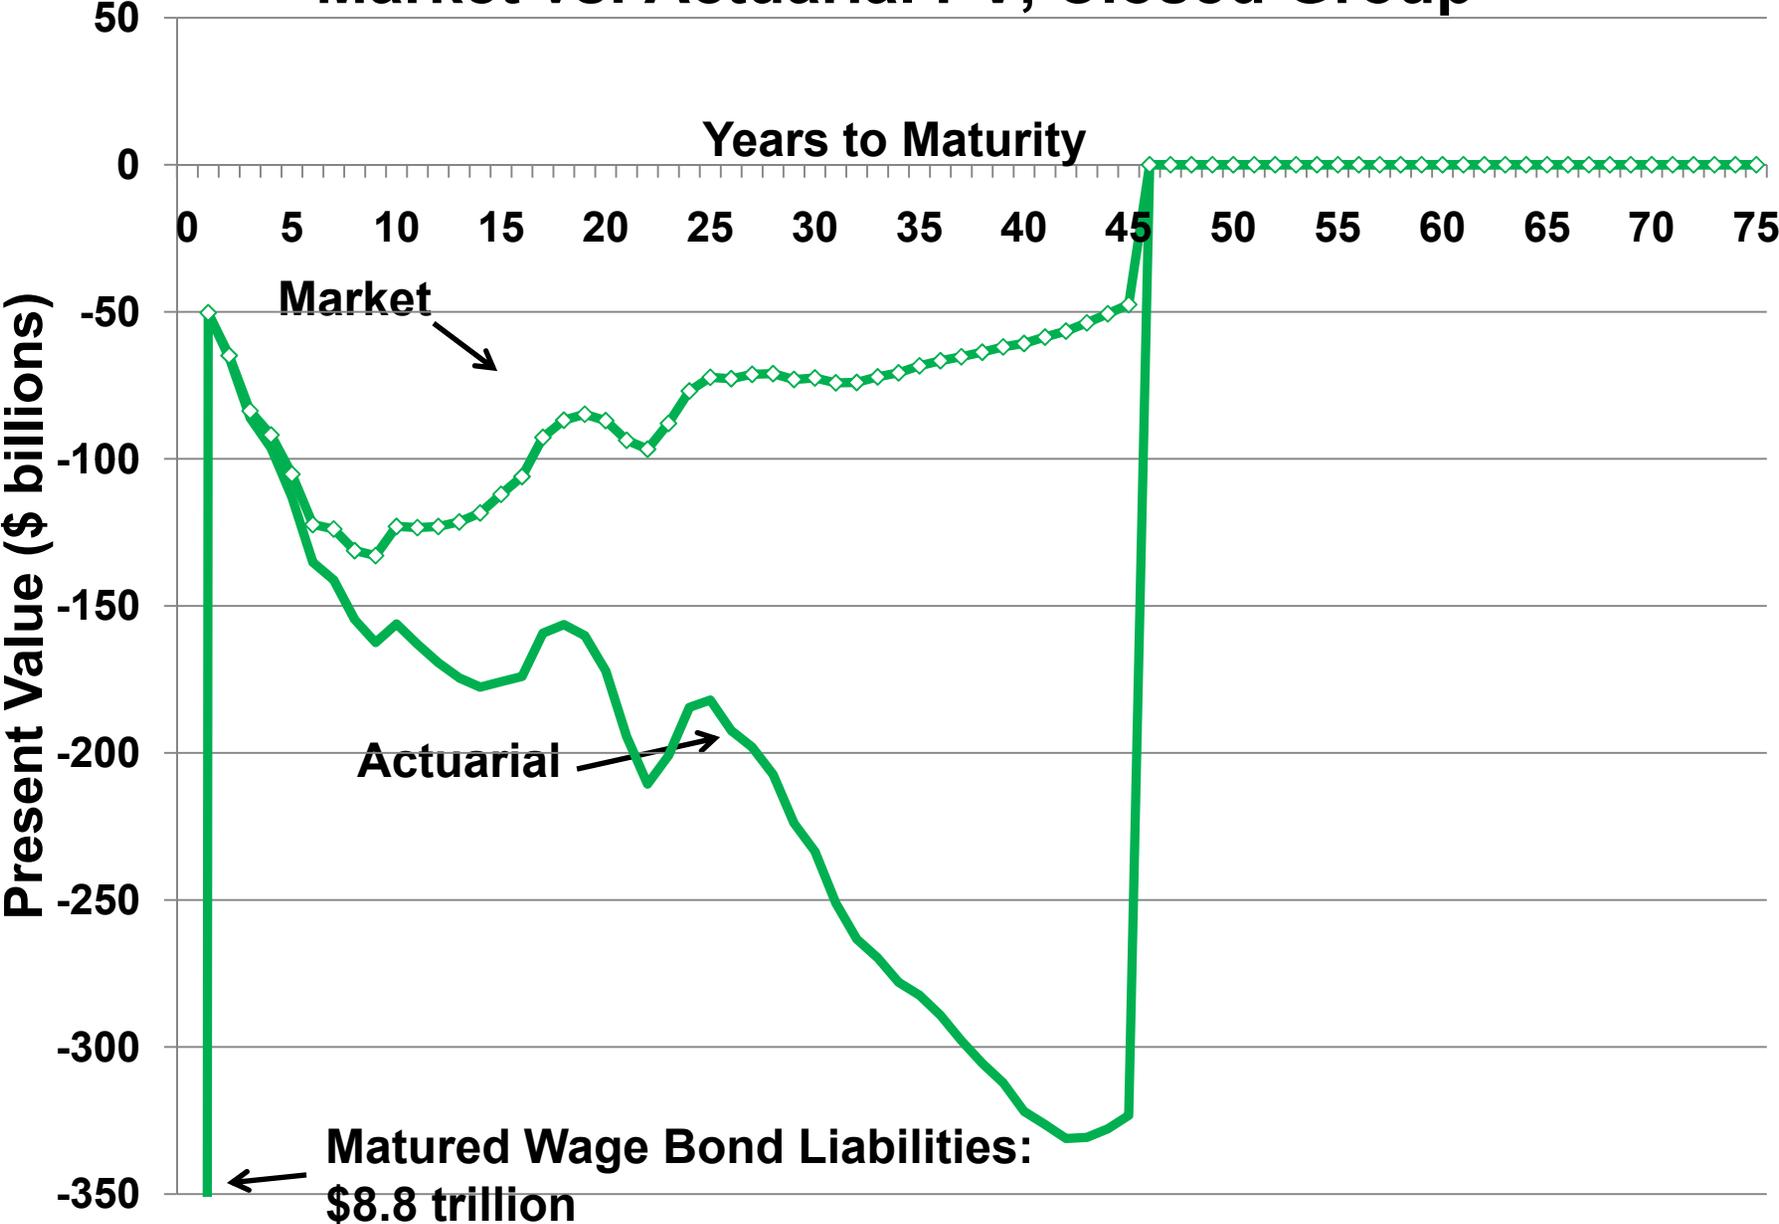


Figure 9b: Net Value of Wage Bonds, Market vs. Actuarial PV, Closed Group



**Figure 9c: Net Value of Wage Bonds,
Market vs. Actuarial PV, 75-Year Open Group**

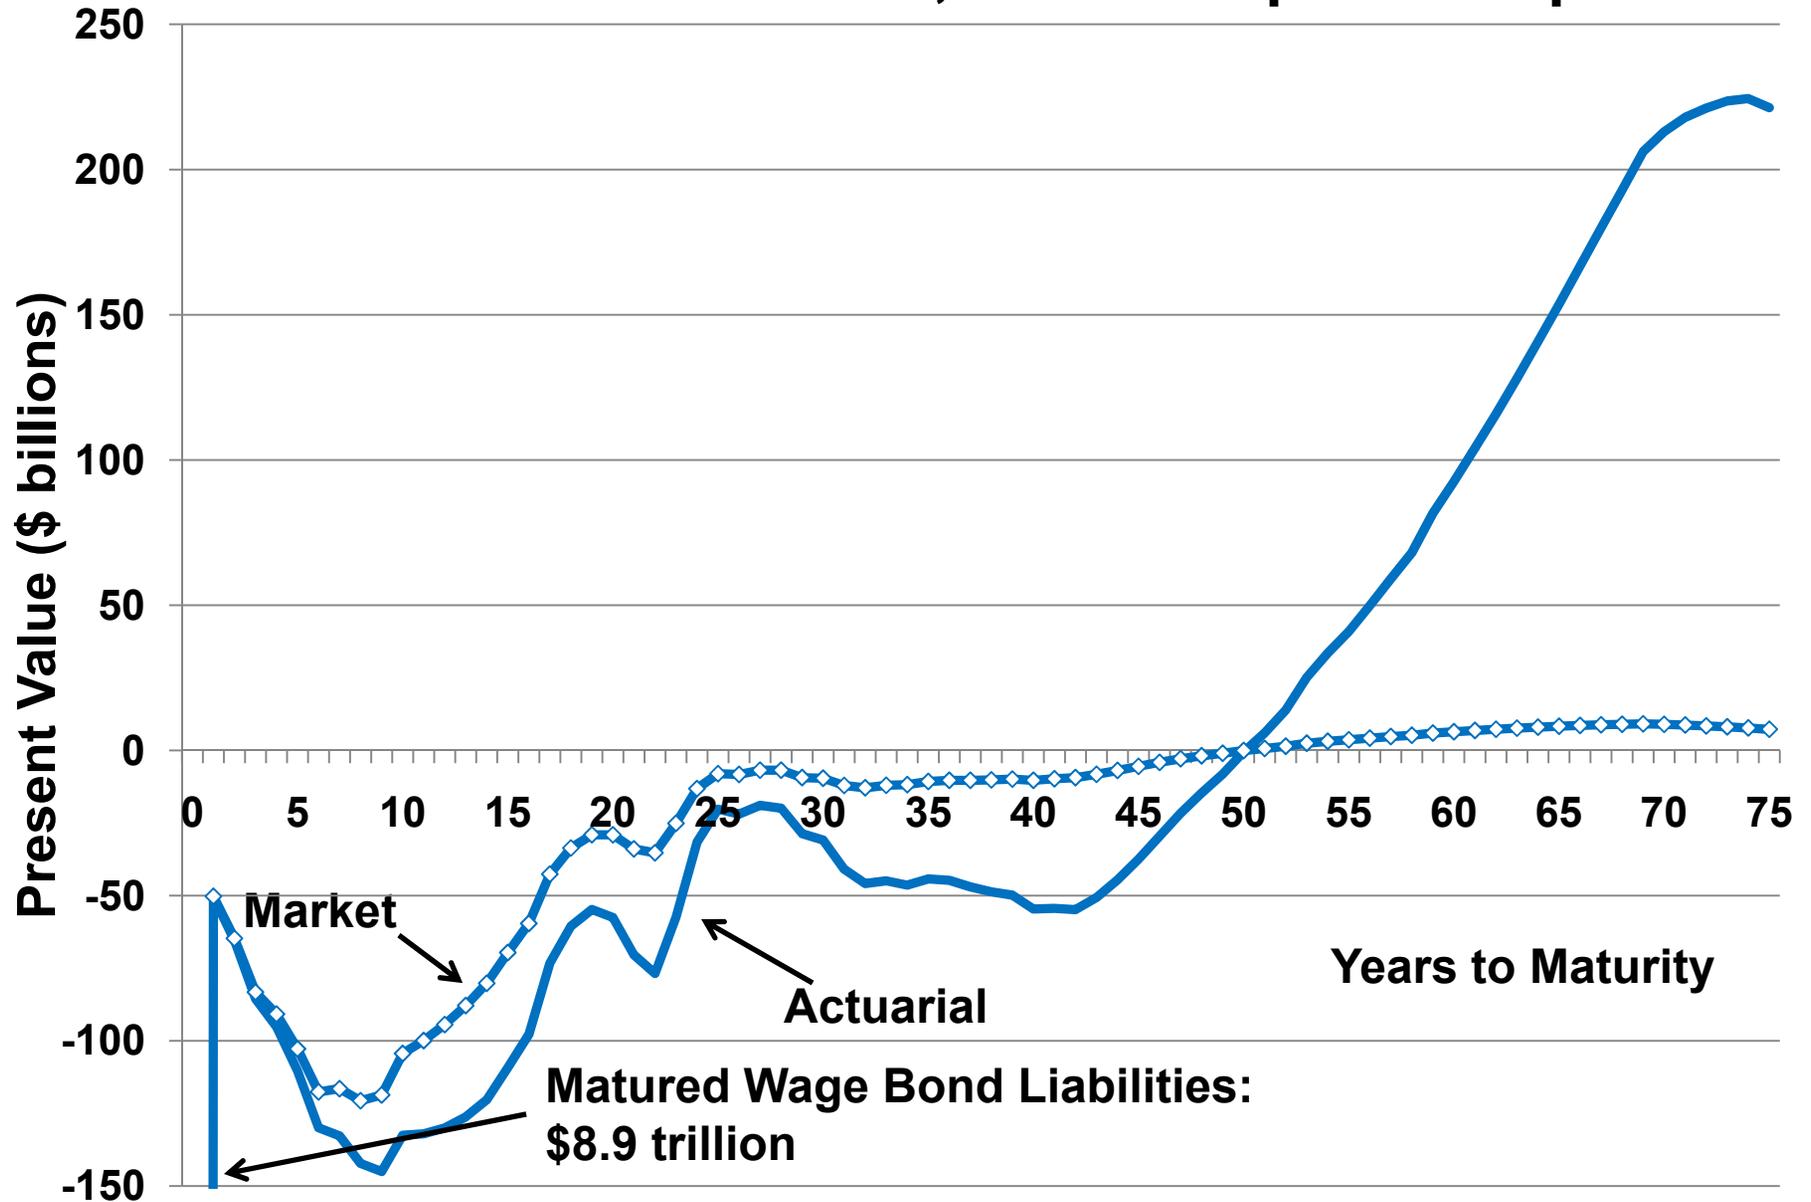
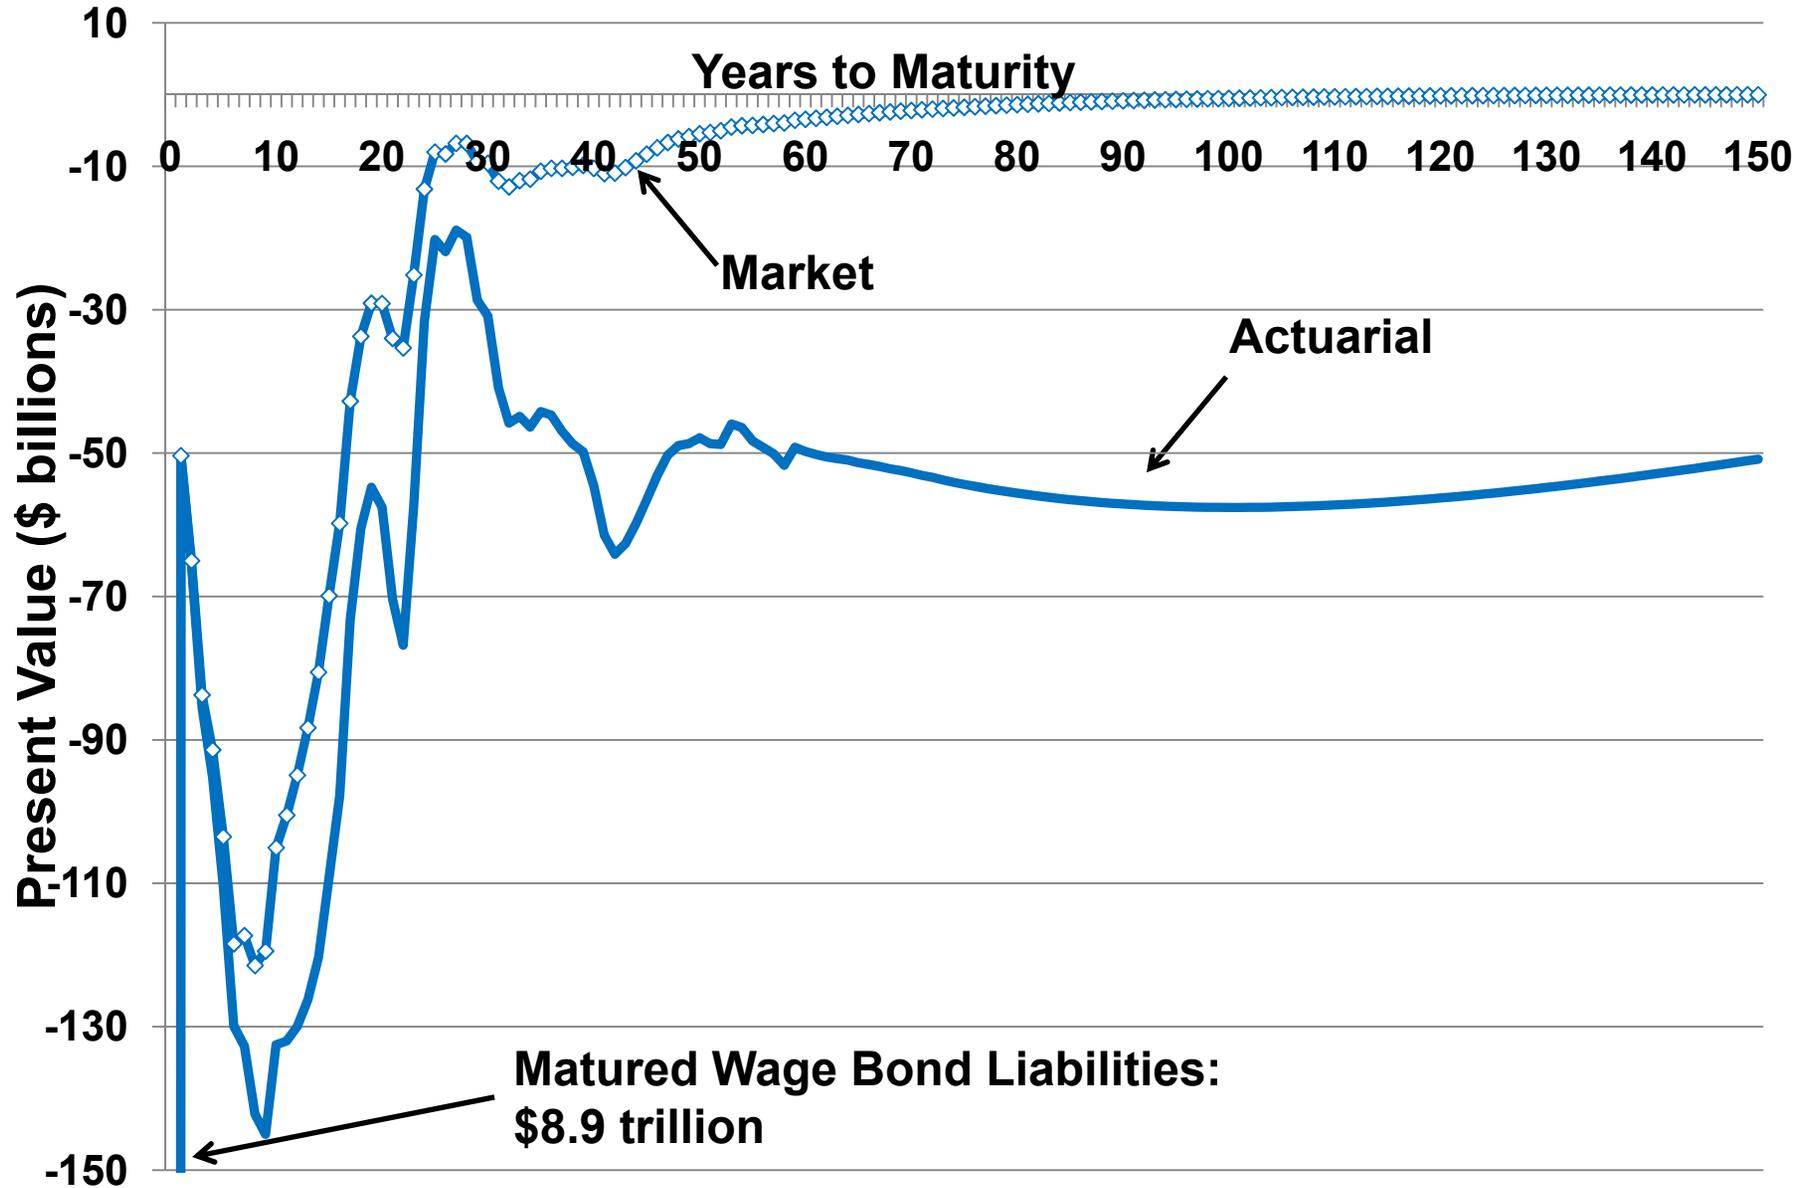


Figure 9d: Net Value of Wage Bonds, Market vs. Actuarial PV, Infinite Horizon Open Group



**Figure 9e: Net Value of Wage Bonds,
Market vs. Actuarial PV, Future Transition Cost**

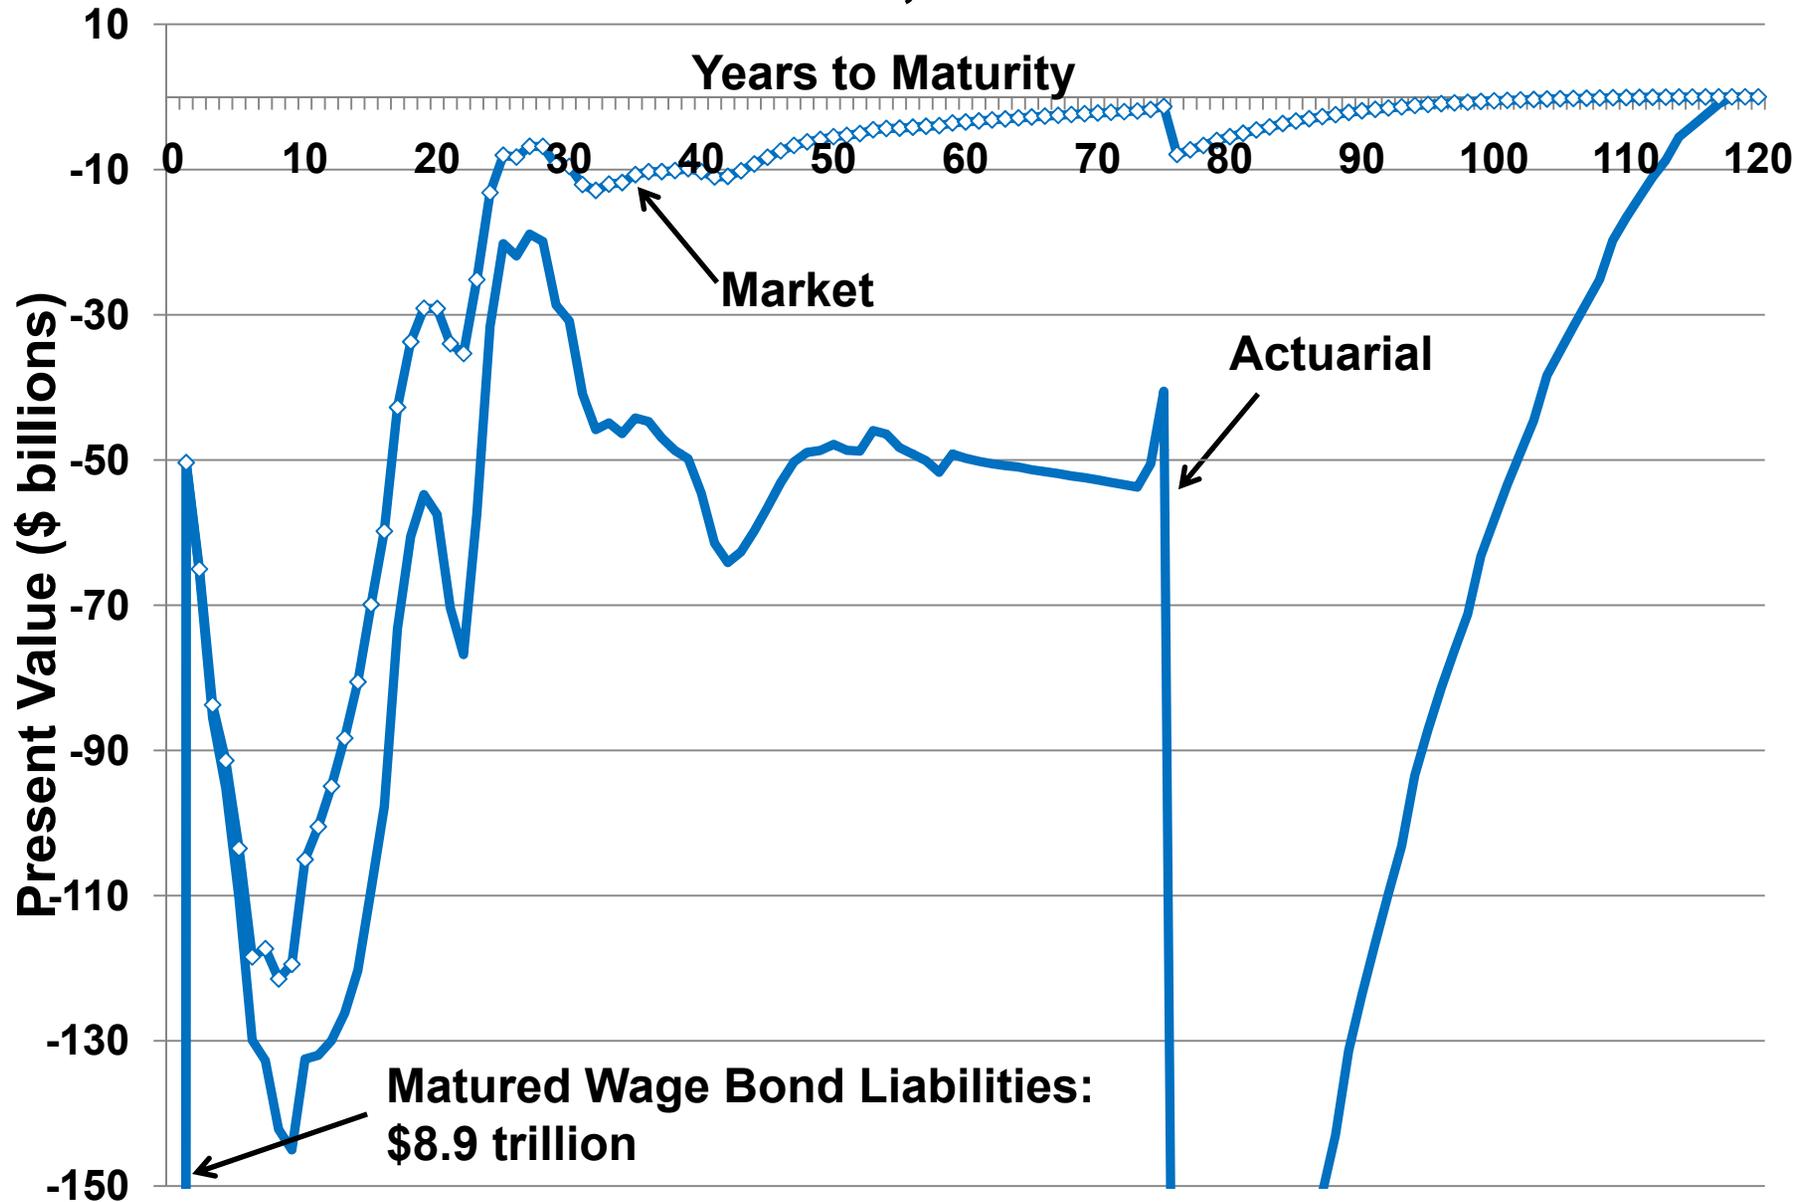


Figure 10: Present values of wage-indexed vs price-indexed cash flows

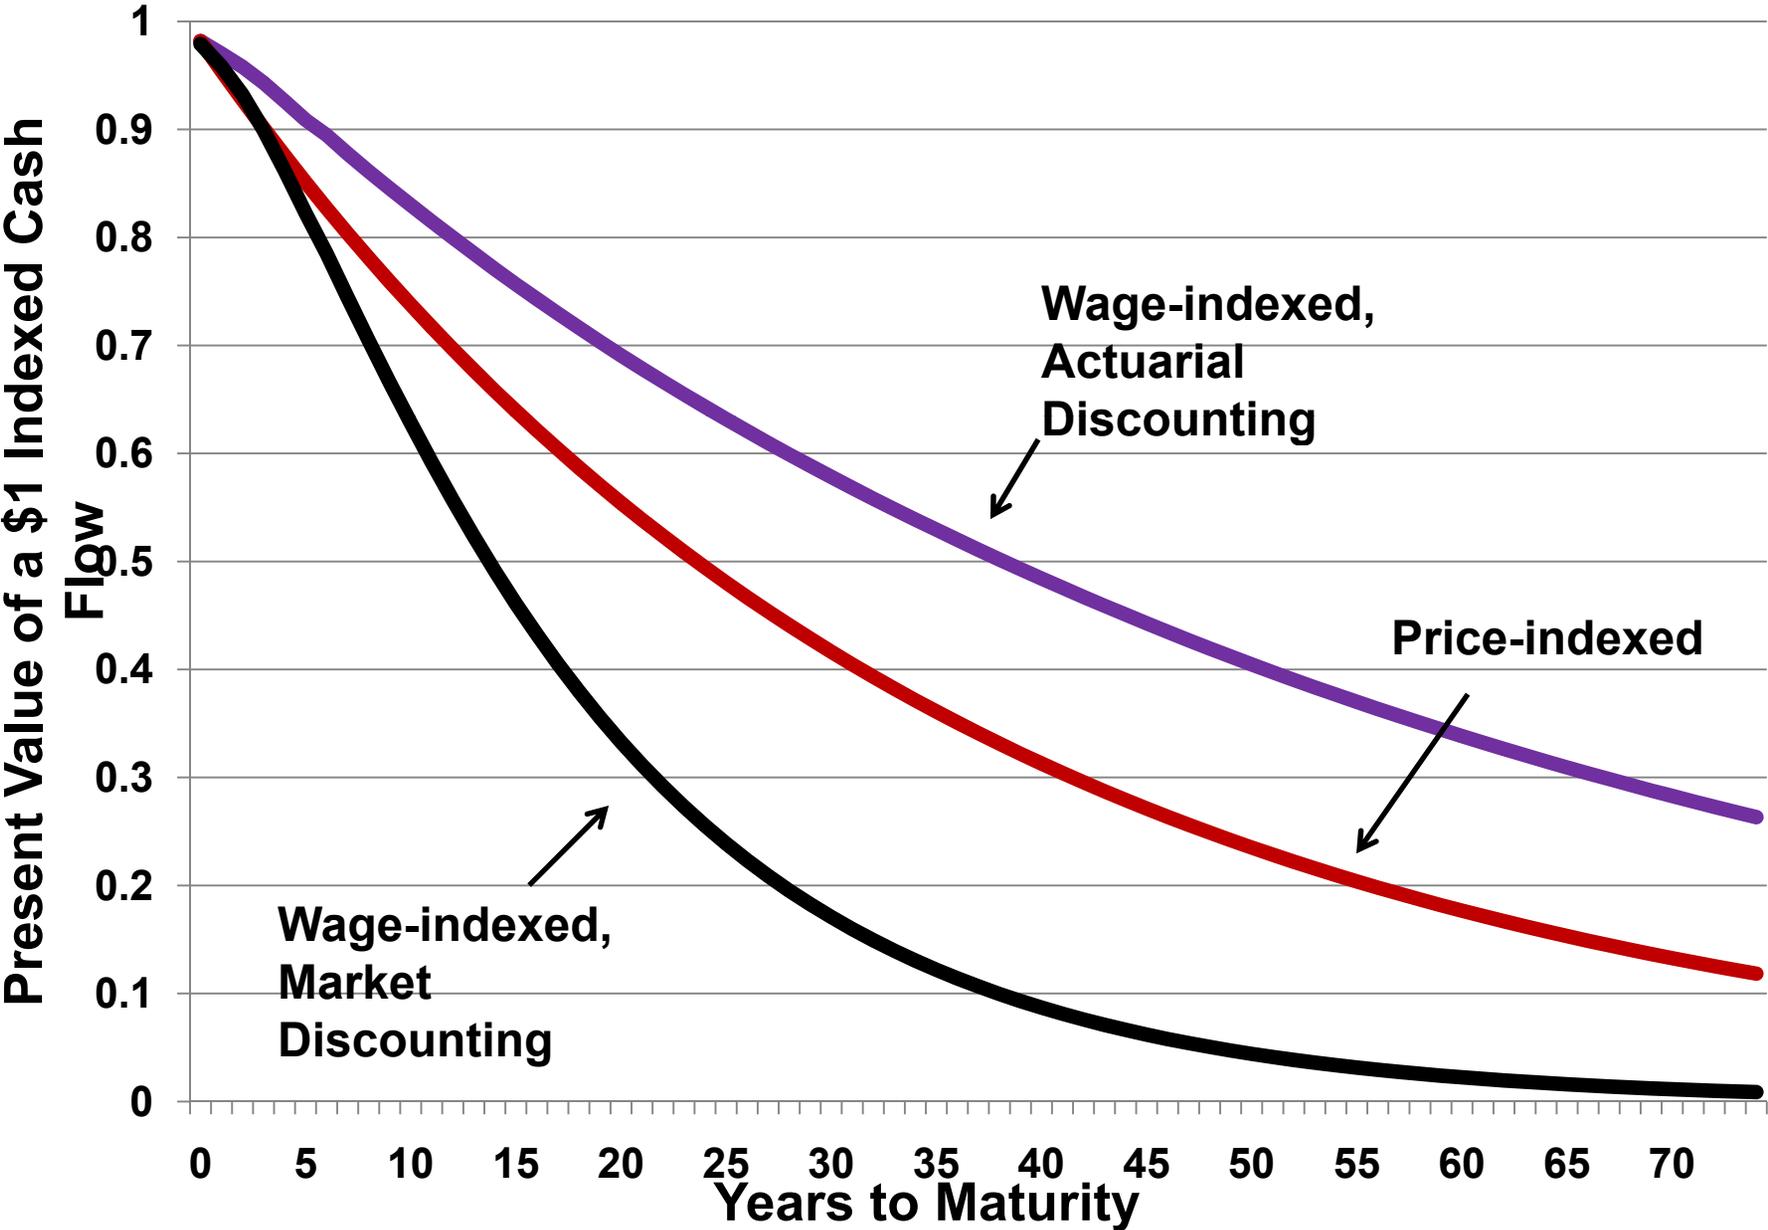


Table 1: Market vs. Actuarial Solvency Estimates (\$ trillion)

	Tax Revenue	Net Benefits	Total Surplus
Immediate Transition Cost			
Actuarial	-	16.7	-14.2
Market	-	14.1	-11.7
Market - Actuarial	-	-2.6	2.6
Market/Actuarial	-	0.85	0.82
Closed Group			
Actuarial	13.8	30.0	-13.7
Market	9.0	19.5	-8.1
Market - Actuarial	-4.8	-10.4	5.6
Market/Actuarial	0.65	0.65	0.59
500-Year Open Group			
Actuarial	44.0	66.4	-19.9
Market	11.7	20.4	-6.4
Market - Actuarial	-32.4	-45.9	13.5
Market/Actuarial	0.26	0.31	0.32

Note: Total surplus = Tax revenue – Net benefits – 2.5 trillion (TF)

Table 2: Market vs. Actuarial Solvency Estimates (\$ trillion)

	Tax Revenue	Net Benefits	Total Surplus
75-Year Open Group			
Actuarial	28.2	35.0	-4.3
Market	11.5	20.0	-6.1
Market - Actuarial	-16.7	-15.0	-1.7
Market/Actuarial	0.41	0.57	1.39
Future Transition Cost			
Actuarial	28.2	43.7	-13.1
Market	11.5	20.4	-6.4
Market - Actuarial	-16.7	-23.3	6.7
Market/Actuarial	0.41	0.47	0.49

Note: Total surplus = Tax revenue – Net benefits – 2.5 trillion (TF)

Table 3: Policy Experiments

	Total Surplus	Tax on 500-Year Payroll	Tax on Relevant Payroll	Benefit Reduction
Immediate Transition Cost				
Actuarial	-14.2	4.0%		85.5%
Market	-11.7	12.4%		82.8%
Closed Group				
Actuarial	-13.7	3.9%	12.3%	45.7%
Market	-8.1	8.6%	11.2%	41.6%
500-Year Open Group				
Actuarial	-19.9	5.6%	5.6%	30.0%
Market	-6.4	6.8%	6.8%	31.2%
75-Year Open Group				
Actuarial	-4.3	1.2%	1.9%	12.4%
Market	-6.1	6.5%	6.5%	30.3%
Future Transition Cost				
Actuarial	-13.1	3.7%	5.7%	29.9%
Market	-6.4	6.8%	6.9%	31.5%

Table 4: Countries With Negative Real Wage Growth

Horizon Years	# of Countries With Negatives	# of Times Negative	% of Times Negative
1	29	256	25%
5	26	174	20%
10	14	100	13%
20	6	46	10%
30	4	20	8%

Note: 29 countries in total