

Market Valuation of Accrued Social Security Benefits

John Geanakoplos and Stephen P. Zeldes

Citation: Geanakoplos, John, and Stephen P. Zeldes. 2010. "Market Valuation of Accrued Social Security Benefits," in *Measuring and Managing Federal Financial Risk*, ed. Deborah Lucas, 213-233. Chicago: University of Chicago Press.

©2010 by the National Bureau of Economic Research

<http://www.press.uchicago.edu/presssite/metadata.epl?isbn=9780226496580>

8.1 Introduction

One measure of the health of the Social Security system is the difference between the present value of Social Security benefits accrued to date and the market value of the Social Security trust fund. This measure, referred to as the *maximum transition cost*, is comparable to the one used to gauge the fundedness of private defined benefit pension plans and provides an estimate of the cost of switching from a primarily pay-as-you-go Social Security system to a fully-funded one.

How should present values be computed for this calculation in light of future uncertainties? We argue that it is important to use market value. Since claims on accrued benefits are not currently traded in financial markets, however, we cannot directly observe a market value. In this chapter, we therefore use a model to estimate what the market price for these claims would be if they were traded.

In valuing such claims, the key issue is properly adjusting for risk. We

John Geanakoplos is the James Tobin Professor of Economics at Yale University and an external professor of the Santa Fe Institute. Stephen P. Zeldes is the Benjamin Rosen Professor of Economics and Finance at Columbia University and a research associate of the National Bureau of Economic Research.

We thank Ryan Chahrour, Ben Marx, Theodore Papageorgiou, and Sami Ragab for research assistance, and Mark Broadie, Deborah Lucas, Kent Smetters, and seminar participants at the Conference on Measuring and Managing Federal Financial Risk (Kellogg School of Management, February 2007) and at Columbia University for helpful comments and suggestions. This research was supported by the US Social Security Administration (SSA) through grant #10-P-98363-1-04 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 authors and do not represent the views of SSA, any agency of the Federal Government, or the NBER.

contend that the traditional actuarial approach—the approach currently used by the Social Security Administration (SSA) in generating its most widely cited numbers—does not adjust appropriately for aggregate risk in future financial flows. In particular, the SSA methodology computes the expected value of aggregate cash flows for each future date and then discounts these at a riskless rate of interest. Instead, we treat aggregate Social Security payments as dividends on a risky asset, and ask what that asset would be worth if it were traded in financial markets. We call the resulting estimate the *market value* of Social Security obligations. Effectively, market valuation incorporates a risk premium that reflects the market risk of the cash flows being discounted. If benefits are risky and this risk is priced by the market, then market value will differ from actuarial estimates.

Why do we believe that market value is the relevant measure of financial status? Let us begin with a simple example. Suppose that a worker's Social Security benefits were always equal to the dividends of one share of a particular stock. It would be sensible to quote the value of those benefits at the market price of the stock. That would, for example, allow the worker to compare the size of his private portfolio, which might hold shares of the same stock, and his Social Security portfolio of benefits. Similarly for the Social Security system as a whole, if all the promised benefits together were identical to 20 percent of the combined European stock market, then one-fifth of European stocks' market capitalization would be a useful guide to understanding the cost of transitioning to a fully funded Social Security system. The market value can also be seen as the amount that the government would need to pay participants in the financial market to accept its obligations or liabilities.

Under the current methodology, however, the SSA would likely report much larger numbers for this worker's promised benefits, because the SSA numbers would ignore the riskiness of the dividends. Historically, total stock returns have been much higher than the riskless rate. This suggests that stock dividends are indeed subject to the kind of uncertainty that leads cash flows to be more heavily discounted by the market. Of course theory, beginning with the capital asset pricing model, also suggests that stock dividends should be discounted by more than the riskless rate.

This example, linking stock market risk to risk in Social Security benefits, is not as far-fetched as it might appear. Benefits are by no means risk free. The US Social Security system is “wage-indexed”; that is, future benefits are tied directly to the economywide average wage index around the year of the worker's statutory retirement age. (We discuss the precise formula later.) We argue that wages and stock prices are linked in the long run, effectively linking Social Security benefits to the performance of the stock market.

Theoretically, a long-run relationship between wages and stocks is natural. If we believe that fifty years from now American businesses will be failing and paying small dividends, we should expect wages to be low by then as

well. Over the long term, countries with high business profits per capita have also paid high wages. Empirically, Benzoni, Collin-Dufresne, and Goldstein (2007) find evidence of cointegration between stocks and wages over a long sample of US data (1927 to 2004), despite the well-known difficulties of identifying such relationships in finite samples. We believe there is already strong evidence for the wage-stock link; our chapter suggests one more reason why studying this relationship further is important.

Real wages and stock market returns do not seem to be contemporaneously correlated, as Goetzmann (2008) and others have pointed out. But it is crucial to realize that a lack of short-run correlation does not imply the absence of a long-run correlation. Consider a simple thought experiment. Suppose that wages (W) and dividends (D) always moved one for one in a geometric random walk, and that at every period investors could forecast dividends one period in advance with certainty, but had no information about the more distant future. Assuming a constant risk-free interest rate and pricing kernel, the price of the stock would then be $P_t = \varphi D_{t+1}$ for some constant φ . Stock market returns $(P_{t+1} + D_{t+1})/P_t = D_{t+2}/D_{t+1} + 1/\varphi$ would be independent of contemporaneous wage growth $W_{t+1}/W_t = D_{t+1}/D_t$, but in the long run stock levels and wage levels would be nearly perfectly correlated.

To take another example, suppose, following Benzoni, Collin-Dufresne, and Goldstein (2007), that dividends follow a geometric random walk and that wages also follow a geometric random walk with an independent fluctuation, but with a drift that depends on the ratio of current dividends to current wages. Once again we would find almost no short-run correlation between wage growth and stock returns, but it is easy to see that a sustained period of high stock dividends and high stock returns would likely foreshadow a period of high wage growth.

In what follows, we assume that wages and dividends follow this process, so that there is a positive long-run correlation between average labor earnings and the stock market. We then use derivative pricing methods standard in the finance literature to compute the market price of individual claims on future benefits, which depend on age and macro state variables. Finally, we aggregate the market value of benefits across all cohorts to arrive at an overall value of accrued benefits and of the maximum transition cost.¹

We find that the market value of accrued Social Security benefits is substantially less than the “actuarial” value, and that the difference is especially large for younger cohorts. Overall, the market value of accrued benefits is only four-fifths of that implied by the actuarial approach. Ignoring retirees

1. In this chapter, we focus on the maximum transition cost measure of financial status. In ongoing work (Geanakoplos and Zeldes, 2009b) we examine alternative open and closed group measures that incorporate future taxes and future accruals.

(for whom the valuations are the same), market value is only 70 percent as large as that implied by the actuarial approach. This implies that the market value of Social Security's unfunded obligations, as measured by the maximum transition cost measure, is significantly less than the actuarial value commonly presented by the SSA.

This difference by itself might change the public's view of the transition cost of the system, and is therefore reason enough to pursue a measure of market value. Recent suggestions by the Federal Accounting Standards Advisory Board to include Social Security obligations on the US balance sheet make the question of their value especially pertinent.

One logical consequence of our approach is that large decreases in the stock market, such as we saw in 2007 and 2008, should significantly decrease the market value of accrued Social Security benefits. The SSA, by contrast, does not seem to have moved its calculations by much.

In work done after the original version of this chapter was written, Blocker, Kotlikoff, and Ross (2008) also attempt a market valuation of outstanding Social Security obligations. They argue for risk adjustments due to (a) the correlation between wage growth and returns on traded assets and (b) the inflation insurance provided by consumer price index (CPI)-indexed benefits. They empirically estimate the correlations between wage growth and traded assets, and they conclude that the market value of Social Security obligations is *greater* than the actuarial value. In contrast, we reach the opposite conclusion, namely that the market value is *less* than the actuarial value.

One reason for this disparity is that in addition to adjusting for risk, Blocker, Kotlikoff, and Ross also change the risk-free rate of interest to what they argue is a more reasonable value. Based on the term structure for Treasury Inflation Protected Securities (TIPS), Blocker, Kotlikoff, and Ross assume a risk-free rate between 1.5 percent and 2 percent, while the SSA projections assume a rate of 2.9 percent for nearly the entire horizon of its projections. To the extent that the SSA uses too high a risk-free rate, SSA will underestimate the present value of accrued benefits, but this would be felt even if Social Security benefits were not at all risky (and thus required no risk adjustment). Blocker, Kotlikoff, and Ross do not disentangle the effects of the adjustment for risk and the change in the risk-free rate, but it appears to us that their choice of a lower risk-free rate is the primary factor driving their results.

It is difficult to ascertain from the Blocker, Kotlikoff, and Ross paper the size or even the direction of the two true risk adjustments that they make. Regarding risk adjustment for wages (see point [a] previously), Blocker and colleagues focus on short-run correlations of wages and stocks; they estimate the correlation using at most a one-period lag and find it to be small. We argue that even though the short-run correlation is close to zero, the long-run correlation is large and positive, which implies that risk adjustment

should be large and should *decrease* the market value today of a claim on future economywide wages.

Regarding the risk adjustment to the value of the inflation-indexed annuity as of the retirement date (see point [b] previously), we agree that some adjustment for inflation insurance may be appropriate (as reflected in the difference between the real return on nominal bonds and the real return on indexed bonds). However, this inflation risk premium is likely much smaller than the 90 to 140 basis point spread used by Blocker, Kotlikoff, and Ross. We assume this premium is zero in our analysis.²

Our chapter is structured as follows. In section 8.2, we describe why we think that market value is the most appropriate measure for estimating Social Security obligations. Section 8.3 describes how our previous work can be used to frame accrued benefits in terms of units of a potentially tradable financial security (a Personal Annuitized Average Wage [PAAW]). Section 8.4 shows how to price this security, incorporating the market price of risk. In section 8.5, we estimate the quantity of PAAWs outstanding by cohort, and in section 8.6 we combine the information in 8.4 and 8.5 to arrive at an estimate of the market value of accrued Social Security benefits. In section 8.7, we consider the robustness of our results to changes in the parameter that determines the strength of the wage-stock link. Section 8.8 concludes.

8.2 The Importance of Market Valuation

Market valuation answers the question: “what payment would financial markets require for taking on the responsibility of paying Social Security benefits?” A market price for Social Security obligations would also provide important information to households, governments, private pension plans, other market participants, and administrators of Social Security. In fact, the 2007 Social Security Technical Panel on Assumptions and Methods (Technical Panel 2007) cited an earlier version of our chapter and recommended that the Trustees of Social Security consider adopting risk-adjusted discount rates.

Finding the market value of Social Security liabilities also implies the ability to hedge them, since valuation and hedging are dual computations. If the Social Security trust fund were someday permitted to diversify out of government bonds, this would provide a valuable guide to determining the optimal portfolio allocation.

It is worth noting that the measure we compute ignores the general equilibrium effects of selling the full quantity of the asset; bringing all Social Security obligations to market at once could well change how the market

2. Note that the measure of financial status that Blocker, Kotlikoff, and Ross (2008) examine (a closed group measure that includes future taxes and future accrued benefits of current workers) differs somewhat from ours, but this cannot explain the difference in results.

values these and other assets. In this respect, our measure is no different than “market capitalization” in the stock market, or measures of aggregate holdings in real estate.

A market price for Social Security obligations will be especially important for improving government accounting. In its annual Financial Report, the US government produces a balance sheet that summarizes the assets and liabilities of the Federal Government. One controversial aspect of the balance sheet is how to account for social insurance programs. In 2006, the Federal Accounting Standards Advisory Board (FASAB) published a preliminary statement on new standards for social insurance accounting (FASAB 2006). The document described two views. The Primary View, held by the majority of the board, would recognize every accrued benefit as a liability of the system.³ Under this view, liabilities should be based on expected benefits “attributable” to earnings to date, using current benefit formulas. In contrast, the Alternative View advocates continuing the current practice of acknowledging only those benefits that are “due and payable” at the time of valuation. Essentially, under the alternative view only current-period benefits not yet paid to beneficiaries (an amount close to zero) would be counted as a liability.

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

In a November 2008 update of the statement (FASAB 2008), FASAB proposed a compromise between these views: accrued benefit “obligations” are to be provided in a note on the federal financial statements, and another measure referred to as the closed group measure (equal to the accrued obligations to date plus future taxes and future accruals of *current* participants) is to be reported as a separate line just below the balance sheet. If the compromise prevails, measures of Social Security’s future obligations will gain prominence in government financial statements, but no new *liabilities* will be recognized on the balance sheet at this time.

Whether or not one wishes to characterize future benefit obligations as “liabilities,” correctly computing their value is essential. It is widely agreed that some measure of the present value of future cash flows should be

3. Accrued benefits would be those earned by fully-insured participants (e.g., Social Security participants who have achieved forty-quarters of covered earnings, the minimum to receive benefits) based on their earnings histories to date.

reported, even if not on the balance sheet. Proper valuation of these risky flows will be essential to the new guidelines' efficacy in accurately portraying the financial status of the Social Security program.

For individuals, a market price for cohort benefits would provide information about the market value of their own benefits, helping them with financial planning decisions regarding saving and asset allocation. A true market price would allow individual households to consider Social Security benefits as any other asset in their portfolio. The cohort-specific estimates in this chapter give some idea of the value of new benefit accruals and how they compare with tax contributions. Workers could compute, for example, a market-based "money's worth" measure such as the ratio of the present value (PV) of benefits to the PV of contributions (for a further description of money's worth measures, see Geanakoplos, Mitchell, and Zeldes [1999]). A market value for benefits would also likely make it more difficult for the government to take them away, enhancing property rights.

Finally, if markets for bonds indexed to Social Security obligations actually develop in the future, buyers and sellers of these new securities would be forced to make the same kind of computations we propose here. If the private sector were permitted to issue these securities, the government could purchase them from the private sector in order to cover a portion of the benefit obligations accrued each year.

8.3 Translating Accrued Benefits into Units of Marketable New Securities (PAAWs)

Under current Social Security rules, workers and employers together contribute 12.4 percent of "covered earnings" (i.e., all labor income up to the earnings cap, equal to \$102,000 in 2008). Upon retirement, workers receive benefits that are linked to their earnings history, and in a particular way, to average earnings in the economy. For each year in the worker's history, earnings are divided by the average economywide wage index from that year, and then multiplied by the average economywide wage index in the computation year (typically age sixty).⁴ Since a worker's benefits depend crucially on average wages in the computation year, they are subject to a type of aggregate risk. In this chapter, we price this risk.

The maximum transition cost is reported annually in a recurring note from the Office of the Actuary (Wade, Schultz, and Goss 2008), and is intended to represent the present value of benefits accrued by current and past workers, net of current trust fund assets. Estimating this measure requires establishing what it means for benefits to be accrued. By definition, accrued benefits

4. In Geanakoplos and Zeldes (2008), we assumed all wages were indexed to age sixty-five wages. Under SSA rules, however, wages after age sixty-two are included at their nominal levels in the formula while wages from earlier years are indexed to economy average wages in the individual's sixtieth year. Thus, aggregate wage risk in a cohort is resolved after year sixty.

can rise, but never fall. In Geanakoplos and Zeldes (2009a), we show that there are many feasible accrual rules and describe two natural rules in detail. For simplicity, we focus here on one of these, “the straight-line” accrual rule, in which accrued benefits to date are defined by setting future wages equal to the worker’s average wage to date and prorating the resulting benefits by a scale factor related to years of work.⁵ This is a relatively conservative accrual rule (in the sense of delaying accrual) and thus tends to decrease the accruals of younger cohorts. Since these are the cohorts for whom the risk adjustment is important, this accrual rule tends to decrease the magnitude of the overall risk adjustment. We show that, even with this accrual rule, the risk adjustment is quite significant.

In Geanakoplos and Zeldes (2009a), we described how to create a system of personal accounts that achieves many of the core goals of supporters of the current system, including risk-sharing and redistribution. We called these “Progressive Personal Accounts.” One step in that process was to show that a personal account system could be structured to exactly reproduce the benefits promised under the current system. This involved the creation of a new financial security, which we named a Personal Annuitized Average Wage security, or PAAW for short. Whether or not Progressive Personal Accounts are adopted, this equivalence means that establishing a price for this theoretical security is sufficient for pricing existing Social Security obligations.

We define a PAAW as a security that pays its owner one inflation-corrected dollar for every year of his life after the year (t_R) in which he hits the statutory retirement age (R), multiplied by the economywide average wage (W_{t_c}) in the computation year (t_c) that he hits age sixty. The PAAWs are tied to specific individuals, indexed by i , through their mortality, the wage index in their cohort’s computation year, W_{t_c} , and the year of the first payout on their security (t_R). In this chapter, we assume all workers retire at sixty-five, fixing the relationship between t_c and t_R . In this context, the notation $\text{PAAW}(i, t_R)$ identifies the relevant information for any PAAW.

Each additional dollar that an individual earns generates additional

5. Specifically, we compute average relative earnings over all years the worker has earnings, up to thirty-five years. If the worker has earnings from more than thirty-five years, we take the average over the thirty-five highest earning years. Average relative earnings are then entered into the current Primary Insurance Amount (PIA) formula, and the result is prorated by $\min\{1, (\text{work years}/35)\}$. For example, if a worker has worked for twenty-five years (equal to $5/7$ of thirty-five years), we average the relative earnings from just these twenty-five years (effectively setting future wages equal to this average), compute the resulting number of PAAWs using the PIA formula, and then multiply the result by $5/7$. Note that this is not identical to the SSA procedure for calculating accruals for their Maximum Transition Cost measure (they average the best $4/5$ of earnings years and scale PIA by $(\text{age}-22)/40$), but the two procedures give similar results. An alternative accrual method, also described in Geanakoplos and Zeldes (2009a), is one we call the “fastest” accrual method, which sets future wages to zero and does not prorate, giving more rapid accruals by adjusting for age before the (progressive) calculation of PIA rather than after. (This is termed “fastest” because no other possibilities exist that have faster accumulation and also satisfy the constraint that accrued benefits will not fall even if future earnings are all zero.) See Jackson (2004) for a further discussion of accrual accounting.

accrued benefits or PAAWs. At any point in time t , an individual's accrued benefits can be summarized completely by the number of PAAWs owned. The present value of accrued benefits is therefore equal to the quantity of accrued PAAWs (known at time t) multiplied by the present value of a PAAW(i, t_R).

The PAAW valuations should differ for individuals in the same age cohort with different mortality probabilities. For example, the longer life expectancies of women means their PAAWs would be more valuable, if they were traded separately. We assume that all members of a birth cohort have the same age profile of survival probabilities.⁶ In the following sections, we examine how to price PAAWs for each cohort, and we then estimate the quantity of PAAWs outstanding and the market value of these PAAWS for each cohort.

8.4 The Price of a PAAW

In Geanakoplos and Zeldes (2009a), we argued that if the Social Security system either required workers to sell a small fraction of their PAAWs or issued extra PAAWS, these securities could be pooled together and sold to financial markets. In this section, we estimate what the market price of these pooled PAAWs would be if they were traded in financial markets. To do so, we develop a valuation model that links the risk in PAAWs to the risk in an asset that is already priced, namely stocks. We compare this value with the value generated from the same model, but ignoring the adjustment for risk. We refer to these respectively as the “market” (or “risk-adjusted”) and “actuarial” (or “unadjusted”) values.⁷

8.4.1 Methodology

The PAAW payouts are tied to average economywide wages in a specific year in the future. They are therefore tied to the macroeconomy and potentially to the stock market. Lucas and Zeldes (2006) show how to value defined benefit (DB) pension liabilities when payouts are tied to future wages of the individual. We apply that approach here, modifying it to take into account the specifics of Social Security benefit rules. One important difference between the two applications is that under private DB pensions, the accrued benefit obligation (ABO) depends only on past labor earnings (and thus requires no risk adjustment), while the projected benefit obligation (PBO) depends on future labor earnings. Due to the wage indexing of Social Security, even the ABO measure of Social Security depends on

6. To the extent that there is a correlation between life expectancy and number of accrued PAAWs, we will underestimate the value of each cohort's accrued PAAWs.

7. A comparison of the risk-adjusted and actuarial values could be used to back out an estimate of the appropriate risk-adjusted discount rate. We pursue this in Geanakoplos and Zeldes (2009b).

future (economywide) labor earnings, and therefore even the ABO measure of Social Security requires an adjustment for salary risk.

The cash flow stream on a PAAW(i, t_R) depends on the economywide average earnings index W_{t_c} at time t_c , the life span of individual i , and the year of retirement t_R . In particular, an individual's retirement benefits are an annuity proportional to the average wage in his sixtieth year. If we define a wage bond as a security that pays an amount equal to the average wage in some future year, then we can decompose the problem of pricing a PAAW into the problem of pricing the wage bond (which requires a model of wage growth), and pricing the annuity (which we assume is independent of wage growth). We proceed in this manner, first pricing the wage bond, then combining our result with a standard valuation for the cohort-specific annuity.

The key issue for pricing the wage bond is the correlation, at different horizons, between aggregate wages and dividends, and thus the value of the stock market. To model this relationship, we use a simplified, discrete-time version of the model used in Benzoni, Collin-Dufresne, and Goldstein (2007). We model the relationship between real variables and assume that inflation does not affect the relationship between real wages and real dividends. We begin with a stationary geometric random walk process for log real dividends (d):

$$(1) \quad d_{t+h} - d_t = h \left(g_d - \frac{\sigma_d^2}{2} \right) + \sigma_d \sqrt{h} z_{d,t+h}.$$

The dividend growth shock, $z_{d,t+h}$, is assumed to be standard normal.⁸

Benzoni and colleagues assume a stationary pricing kernel with a constant price of risk, λ . This implies a constant price-dividend ratio, and therefore a constant dividend yield, δ .⁹ Because the stock price is proportional to current period dividends, it too will follow a geometric random walk with the growth in the stock price exactly equal to the growth in dividends. The total real stock return (r^s) thus equals the dividend yield *plus* the growth in real dividends:

$$(2) \quad r_{t+h}^s = h\delta + (d_{t+h} - d_t) = h \left(g_d + \delta - \frac{\sigma_d^2}{2} \right) + \sigma_d \sqrt{h} z_{d,t+h}.$$

8. 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 .

9. We can see this from the present value relationship,

$$P_0 = E \sum_{t=0}^{\infty} (1 + r + \lambda \sigma_d)^{-t} D_t = D_0 \sum_{t=0}^{\infty} (1 + r + \lambda \sigma_d)^{-t} (1 + g_d)^t.$$

Computing the sum, we have

$$\frac{P_0}{D_0} = \frac{1}{1 - (1 + g_d)/(1 + r + \lambda \sigma_d)} \approx \frac{1}{r + \lambda \sigma_d - g_d}.$$

In continuous time, the last statement is an exact equality.

Note that equation (2) implies the counterfactual result that stock returns and dividend growth have the same volatility.

Next, we describe the process for log real wages (w_t), in which log wage growth is a function of (a) a deterministic wage growth, or “drift”, parameter; (b) the current-period deviation from the long-term average wage-dividend ratio; and (c) an independent and identically distributed (i.i.d.) wage growth shock:

$$(3) \quad w_{t+h} - w_t = h \left(g_w - \frac{\sigma_w^2}{2} \right) - h\kappa(w_t - d_t - \overline{wd}) + \sigma_w \sqrt{h} z_{w,t+h}.$$

In this model, wage growth tends to correct deviations in the wage-dividend ratio from its long term level, \overline{wd} . The parameter κ determines the rate at which the wage-dividend ratio “error corrects.”

As a baseline calibration, we choose parameters that are consistent with the 2008 Trustees Report intermediate cost assumptions. As discussed before, Blocker, Kotlikoff, and Ross argue that this is not the most reasonable parameterization. In order to emphasize the role of *risk-correction*, however, we believe this is the best starting point. Accordingly, the real risk-free rate, r , is set to 2.9 percent and average real wage growth, g_w , to 1.1 percent. In addition, we choose the dividend yield, δ , in order to match the empirical equity premium, which we estimate to be 5.1 percent annually over the period from 1959 through the first half of 2008 and we set g_d to 1.1 percent (equal to g_w).¹⁰ Note that this implies a counterfactually large dividend yield, δ , of 6.9 percent = 5.1 percent – 1.1 percent + 2.9 percent. Finally, we set σ_d (the standard deviation of stock returns and dividend growth), equal to 12 percent, based on the volatility of real stock returns in our sample.¹¹

From the perspective of this chapter, the most important parameter calibration is our choice of κ . Benzoni, Collin-Dufresne, and Goldstein (2007) estimate κ to be between .05 and .2, and take 0.15 as their baseline value, which we follow in this chapter. We also examine the robustness of our results to different values of κ .

Following Lucas and Zeldes (2006), we assume that all risk not captured by the relationship between wages and stocks would be priced by the market at zero, and we use risk-neutral Monte Carlo derivative pricing techniques (as in Cox, Ross, and Rubenstein 1979) to price a wage bond as a derivative

10. Our estimate of the equity premium is equal to the (arithmetic) average of the monthly return on the S&P 500 index minus the average interest rate on three-month T-bills.

11. Benzoni, Collin-Dufresne, and Goldstein (2007) assume an equity premium of 6 percent and use the parameter configuration $g_d = 1.8$ percent, $r = 1$ percent, and $\sigma_d = 16$ percent. We have selected g_d and r to best match the assumptions underlying the SSA actuarial estimates, even though these choices may be controversial. Because of Jensen’s terms in the wage process, however, $E(W(t+n)/W(t))^{1/n}$ is increasing over time. Thus, although we match the actuarial projection of wage growth year-over-year, cumulative wage growth increases to an annualized rate of 1.6 percent at the forty-year horizon. In levels, expected wages are about 20 percent higher at this horizon than they are under the SSA expected growth assumptions.

on the stock market. This entails generating a set of hypothetical “risk-neutral” probabilities on the set of possible returns for stocks such that, under those probabilities, the expected return would equal the risk-free rate. In our simple model, this “risk-neutral” distribution for stock returns is normal with a mean equal to the risk-free rate and the standard deviation equal to its original empirical value.

We use Monte Carlo techniques to simulate stock returns and wages using the risk-neutral probabilities. We generate 200,000 replications of the wage and dividend process, each forty-five years in length, and take averages over the realizations. Our estimate of the “risk-adjusted” price of a year- t wage bond is equal to the average value of the simulated wage at year t , using risk-neutral probabilities, discounted at the risk-free rate.

We use the wage bond price to compute the current market value of a PAAW. A PAAW for this worker promises payments proportional to the age sixty average wage, starting in the retirement year, which we assume to be age sixty-five. To compute annuity prices, we use the cohort life tables from Bell and Miller (2002) and assume that all individuals of the same age face the same conditional survival probabilities¹² (i.e., that there is no heterogeneity or private information about these probabilities). We also assume that the market price of aggregate longevity risk and inflation risk are each zero.

As a concrete example of how we compute PAAW prices, consider the cohort of age fifty, which reaches age sixty in 2015, ten years from our valuation date. We compute the market (risk-adjusted) value in 2005 of the 2015 wage bond to be 0.658 current wage units. The age sixty value of a one dollar perpetual real annuity starting at age sixty-five, valued using cohort-specific mortality and a risk-free rate, is \$10.88. Finally, conditional on being fifty years old in 2005, there is a 92.3 percent chance of reaching age sixty, the year we value the annuity. Therefore, the 2005 market value of a PAAW for this cohort is $(10.88) \cdot (0.658) \cdot (0.923) = 6.60$ current wage units. Multiplying by the current value of the average wage gives the market value of a PAAW, measured in dollars.

8.4.2 Actuarial Value

The standard actuarial approach for computing present value makes no adjustment for risk; that is, it computes the expected value of the cash flows and discounts at the risk-free rate.¹³ To estimate the “non-risk-adjusted” or actuarial price of a wage bond, we use the same model described before, but generate a set of wage and dividend realizations that are based on the true

12. For the calculations presented, we used the survival probabilities for males born in 1980. Using sex-specific survival probabilities increases our measure of accrued benefits by about 7 percent (since women typically live longer than men). The risk-adjustment, however, is only negligibly affected.

13. Note that if all individuals in the economy were risk-neutral, no adjustment for risk would be necessary, and the actuarial and market approaches would yield identical results.

probabilities, and then discount the average value of the simulated wage at the risk-free rate. We use the resulting wage bond price to compute the actuarial price of a PAAW. In the example above, the actuarial value in 2005 of a 2015 wage bond is 0.830 wage units (versus a market value of 0.658). The resulting actuarial value in 2005 of a 2015 PAAW is $(10.88) \times (0.830) \times (0.923) = 8.34$ current wage units (versus a market value of 6.60).

8.4.3 Results

Figure 8.1 compares the actuarial and market prices of the wage bonds. The risk adjustment causes the market price to be everywhere lower than the actuarial price. In addition, the difference grows over time, since wages further out are more risky and subject to a larger adjustment.¹⁴

Figure 8.2 compares the actuarial and market prices of PAAWs. Figure 8.3 shows the ratio of market (risk-corrected) to actuarial PAAW prices for each cohort. For cohorts that have already surpassed the computation age (sixty), the risk-adjustment has no impact on the valuation. This occurs because aggregate wages are the only source of priced risk in our model, and cohort benefits depend on aggregate wages in the year it turns sixty. For younger cohorts, however, there is a significant difference between the two methods. For cohorts under age forty, the market measure is less than half of the actuarial valuation. For the youngest cohorts we consider (age twenty in 2005), risk-adjusted accruals are worth less than 20 percent of their value under the standard approach.

8.5 The Quantity of PAAWs Outstanding

In this section, we estimate the stream of future benefits that have been accrued by each cohort based on contributions to date. As pointed out previously, these can be neatly described with a single summary statistic: the number of PAAWs accrued by the cohort.

To construct accrual, we use data from the Continuous Work History Sample (CWHS), a 1 percent sample of workers and beneficiaries.¹⁵ The key feature of this data set, for our purposes, is that it includes individual-specific earnings histories.¹⁶ We compute accrued benefits for both current

14. Both prices decrease with the horizon, reflecting the fact that the risk-free rate is greater than average wage growth. In addition, both prices are slightly less than one in the initial 2005 period due to our assumption that cash flows occur at the end of each period and are discounted back to the beginning of the period.

15. We are grateful to Jae Song and Wojciech Kopczuk for providing us with summary statistics from the CWHS.

16. Earnings occurring before 1951 are treated differently in this data set and are typically available only as single entry summing all earning from 1950 and earlier. We ignore these earnings entirely, meaning we slightly underestimate benefits for the oldest cohorts we consider. Because the benefit formula allows workers to exclude low earnings years, typically early years in a worker's history, our underestimate should be very small.




Fig. 8.1 Wage bond prices




Fig. 8.2 Price-per-PAAW

and former workers (including retirees). For retirees this simply entails averaging the thirty-five years of highest relative earnings and entering this average into the Primary Insurance Amount (PIA) formula (redefined to be in units of future economywide wages). For workers who have not already retired, we use the straight-line accrual formula described before to compute PAAW accruals based on worker earnings histories to date. Because our data




Fig. 8.3 PAAW price ratios

set has no information on spousal earnings or status, our results ignore any potential spousal or survivor benefits. The quantity of PAAWs accrued to date by a cohort is equal to the sum of the PAAWs accrued to date by all individuals in the cohort.

8.5.1 Estimates of PAAW Quantities by Cohort

Figure 8.4 shows our estimate of PAAWs earned through 2004 for cohorts born between 1910 and 1986 (ages nineteen through ninety-five in 2005). The hump shape in quantities reflects three key features of benefit accruals and Social Security demographics: (a) younger cohorts have shorter work histories and thus have accrued fewer benefits; (b) the middle-aged cohorts are large and have already accrued most of their benefits; and (c) older cohorts have fewer members because of mortality (for example, in 2005 there were 3.6 million living individuals aged fifty-five but only 2.3 million aged sixty-five and 1.7 million aged seventy-five).

8.6 The Market Value of Accrued Benefits

Once we have computed the price of a PAAW for each cohort and the quantity of PAAWs outstanding for each cohort, estimating the market value of accrued benefits simply involves multiplying the two and summing across cohorts. Figure 8.5 compares the risk-adjusted and the actuarial valuations by cohort. As with the wage bond prices in figure 8.1, the risk-adjustment reduces the value of the liability for all of the nonretired cohorts.




Fig. 8.4 Quantity of accrued PAAWs




Fig. 8.5 Value of accrued PAAWs

Differences across cohorts of the adjustment suggest that risk-correction should be a key consideration in evaluating the “fairness” of proposals to reform Social Security.

Table 8.1 sums accrued benefits across cohorts for an estimate of the total value of accrued benefits. We present two estimates: an actuarial valuation and a risk-adjusted valuation. Our estimate of total accrued benefits, based on the actuarial valuation methodology, is just under \$13 trillion. Adjusting

Table 8.1 Present value of accrued Social Security benefits under alternative valuation methods

	Total value (billions)	Under 60	Over 60
Actuarial (unadjusted)	12,977	8,572	4,405
Market (risk-adjusted)	10,451	6,046	4,405
Market/actuarial	0.81	0.71	1.00

Note: 2006 Office of the Actuary (OACT) Actuarial Note estimate of Max. Trans. Cost + Jan 1st 2006 Trust Fund balance, adjusted to include “own-history” benefits only, equals 12.2 tril.

the Office of the Actuary’s own 2005 estimate of accrued benefits for comparability gives a value of \$12.2 trillion.¹⁷ Given our lack of information about benefits other than basic retirement benefits paid to primary beneficiaries, our estimate of accruals without risk adjustment comes remarkably close to SSA figures.¹⁸

We estimate a market value for the same liability of \$10.5 trillion, only 81 percent of the actuarial value.¹⁹ This difference in valuation comes entirely from the risk-correction; all other features of the pricing model are held

17. Our estimate from CWHS data includes only “own-history” accruals; that is, it excludes spousal and survivor benefits. To obtain a comparable estimate from SSA publications we start with the January 1, 2006 value of the Maximum Transition Cost of \$15.8 trillion, which is the present value of accruals less the amount of the Social Security Trust Fund (Wade, Schultz, and Goss 2008). To this we add back the December 31, 2005 value of the OASDI Trust Fund of \$1.86 trillion (Social Security Administration 2008). We then multiply this sum by the percentage of benefits paid to retired workers based on their own earnings history, which was roughly 70 percent in 2005 (Social Security Administration 2006). To make this adjustment, we assume that the proportion of benefits going to disability and survivors is constant across cohorts and over time. This implies that these programs represent a constant proportion of accrued benefits as well.

18. In principle our actuarial estimate should exactly match the adjusted SSA figure. Differences may arise for at least three reasons: (a) Our limited information does not allow us to perfectly adjust SSA figures derived from micro models. To make this adjustment, we make the simplifying assumption that the proportion of benefits going to spouses, survivors, and disabled beneficiaries is constant across cohorts and over time. (b) The “straight-line” accrual formula we use is slightly different than the one used by SSA to compute the maximum transition cost (MTC) measure, principally because SSA excludes some years of low earnings in estimating PIA, even for workers who have yet to reach thirty-five years of earnings, while we do not (see footnote 5). (c) Expected long-term growth in wages differs from SSA projections, as described in footnote 11.

19. This differs from an earlier (2007) draft of this chapter for three reasons. First, in this version we have linked retirement benefits to wages at age sixty (as opposed to age sixty-five in the earlier draft), effectively removing five years of risk from every cohort. This is appropriate because, as noted earlier, the age sixty wage index is used in computing benefits. Second, in this version, we use the straight-line method of accrual, instead of the “fastest” method used in the earlier draft. We choose this because it more closely matches the measure used by the Office of the Actuary to compute the maximum transition cost estimates. It implies lower current accruals for nonretired workers—those for whom the risk adjustment matters. Under fastest accrual, the corresponding adjustment is 22 percent. Finally, in this draft we are using revised estimates from the 2005 CWHS, whereas in the previous version we used two sources: the 2004 CWHS and a set of OASDI benefit expenditure projections provided by the SSA Office of the Actuary.

constant in generating the figures. This suggests that the standard approach of discounting expected future benefits by the risk-free rate is significantly overstating the size of accrued benefits. Appropriately correcting for risk to aggregate wage growth reduces our measure of Social Security benefits obligations by nearly 20 percent. Subtracting the end of 2005 value of the Old-Age and Survivors Insurance (OASI) trust fund (1.66 trillion) from both measures indicates that the market value estimate of the maximum transition cost measure of Social Security's financial status is only 78 percent as large as the actuarial value, suggesting a healthier system (in the sense of ease of transition to an alternative system) than found using traditional actuarial methods.

Table 8.1 also breaks down the liability for cohorts below age sixty, and those sixty and above. Age sixty is key because that is the year by which the wage risk to benefits is resolved. For the sixty and over group, the actuarial and risk-adjusted estimates are identical, and the aggregate numbers reflect this. When we examine the pre-sixty-year-old group alone, however, we see significantly larger differences between the actuarial and risk-adjusted estimates: correcting for risk reduces our measure of Social Security benefits obligations for those under sixty by nearly 30 percent.

8.7 Robustness

The parameter κ plays a key role in our analysis because it governs the strength of the link between wages and the stock market. Our baseline calibration follows Benzoni, Collin-Dufresne, and Goldstein (2007) in setting this parameter to .15. However, because of the difficulty in estimating such cointegrating relationships, it is informative to examine the sensitivity of our results to this parameter. To do this, we perform the same simulation with a high (.25) and a low (.05) value for κ . Figure 8.6 shows the ratio of the risk-adjusted price to the actuarial price for PAAWs under the alternative calibrations.

First, we find, not surprisingly, that the importance of risk correction varies directly with κ : higher κ implies that wage growth is more "exposed" to stock market risk and increases the size of the risk adjustment.

In addition, we see in figure 8.6 that the size of the risk correction varies in a nonlinear way with κ . For all cohorts, increasing κ from a low value of .05 to our baseline value of .15 has a large effect on the ratio of market to actuarial value, whereas further increasing κ from the baseline to a value of .25 has a much smaller effect.

Finally, the impact of varying κ differs across cohorts. Define the risk adjustment as the distance as measured down from the dashed line. The percentage change in this risk adjustment in response to changing κ is lower for the older cohorts than it is for the younger cohorts. Consider the fifty-year-old cohort as an example. The adjustment represents under 1 percent




Fig. 8.6 PAAW price ratios: Robustness

of the actuarial value under the “low κ ” parameterization, but 27 percent of the actuarial value under the “high κ ” parameter choice. In contrast, for the twenty-year-old cohort, the adjustment is large even for low κ , and raising κ results in a much smaller percentage increase in the adjustment than it did for the fifty-year-old cohort. This pattern is natural; in our model, the long-run correlation between wages and the stock market is 1 for any κ greater than 0, even a small value. Thus, the risk adjustment for benefits far in the future will be (essentially) independent of the parameter κ . On the other hand, the shorter-run correlation between wages and the stock market is highly dependent on κ , so that the risk adjustment of the benefits of workers closer to retirement is much more sensitive to the value of κ .

Table 8.2 aggregates the results across cohorts and examines how they change as κ varies. Increasing κ from the baseline of .15 to .25 increases the risk correction by only 4 percentage points (from 19 percent to 23 percent). On the other hand, lowering κ from .15 to .05 decreases the risk adjustment by 8 percentage points (from 19 percent to 11 percent), a much larger amount. The risk adjustment remains important, however, even with this weak link between wages and stock prices.

8.8 Conclusions, Policy Implications, and Future Research

We have argued that market value is the appropriate way to measure both the assets and the liabilities of the Social Security system. Market value calculations adjust for risk and differ in important ways from the standard actuarial approach that discounts expected cash flows with a risk-free rate.

Table 8.2 Market/actuarial ratio: Robustness to cointegration parameter

	Total	Under 60	Over 60
Low ($\kappa = .05$)	0.89	0.83	1.00
Baseline ($\kappa = .15$)	0.81	0.71	1.00
High ($\kappa = .25$)	0.77	0.66	1.00

We estimate that adjusting for risk reduces the present value of accrued benefits of the entire system by about 20 percent and of workers under age sixty by about 30 percent.

In ongoing work (Geanakoplos and Zeldes 2009b), we extend this approach to consider other measures of Social Security's financial status, including open group measures that incorporate both future Social Security contributions and the corresponding future accruals. Because future tax contributions are proportional to wages (up to the earnings cap), they are subject to a similar risk correction. For the measure we study here, where only future benefit flows must be valued, the direction of the risk adjustment effect is unambiguous; Social Security obligations are worth less under market valuation. Once we consider adjusting both the assets (future taxes) and the liabilities of Social Security (including future accruals), the picture becomes significantly more complicated, and preliminary results suggest that the market value of open group measures shows a larger deficit than the actuarial value.

References

- Bell, F. C., and M. L. Miller. 2002. Life tables for the United States Social Security Area 1900–2100. Social Security Administration, Actuarial Study no. 116.
- 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* 62 (5): 2123–67.
- Blocker, A. W., L. J. Kotlikoff, and S. A. Ross. 2008. The true cost of Social Security. NBER Working Paper no. 14427. Cambridge, MA: National Bureau of Economic Research, October.
- Cox, J. C., S. A. Ross, and M. Rubenstein. 1979. Option pricing: A simplified approach. *Journal of Financial Economics* 7 (September): 229–63.
- Federal Accounting Standards Advisory Board. 2006. Accounting for Social Insurance, revised—Statement of Federal Financial Accounting Standards. Available at: http://www.fasab.gov/pdffiles/socialinsurance_pv.pdf, (October).
- . 2008. Accounting for Social Insurance, revised. Available at: http://www.fasab.gov/pdffiles/socialins_exposurefinal.pdf.
- Geanakoplos, J., O. Mitchell, and S. P. Zeldes. 1999. Social Security money's worth. In *Prospects for Social Security reform*, ed. O. S. Mitchell, R. J. Meyers, and

- H. Young, 79–151. Philadelphia: Pension Research Council, University of Pennsylvania Press.
- Geanakoplos, J., and S. P. Zeldes. 2009a. Reforming Social Security with progressive personal accounts. In *Social Security policy in a changing environment*, ed. J. Brown, J. Liebman, and D. A. Wise, 73–128. Chicago: University of Chicago Press.
- . 2009b. The market value of Social Security. Columbia University Graduate School of Business. Unpublished Manuscript.
- Goetzmann, W. N. 2008. More Social Security, not less. *Journal of Portfolio Management* 35 (1): 115–23.
- Jackson, H. 2004. Accounting for Social Security and its reform. *Harvard Journal on Legislation* 41 (1): 59–159.
- Lucas, D., and S. P. Zeldes. 2006. Valuing and hedging defined benefit pension obligations: The role of stocks revisited. Columbia University Graduate School of Business. Unpublished Manuscript, September.
- Social Security Administration. 2006. *Annual statistical supplement to the Social Security bulletin*, Table 5.A1. Washington, DC: Office of Policy. Available at: <http://www.socialsecurity.gov/policy/docs/statcomps/supplement/2006/5a.html>.
- . 2008. *The 2008 annual report of the board of trustees of the Federal old-age and survivors insurance and federal Disability Insurance trust funds*. Table II.C.2. Washington, DC: Office of Policy. Available at: <http://www.socialsecurity.gov/OACT/TR/TR08/trLOT.html>.
- Technical Panel. 2007. *Report on assumptions and methods*. Report to the Social Security Advisory Board, October. Available at: http://www.ssab.gov/documents/2007_TPAM_REPORT_FINAL_copy.PDF.
- Wade, A., J. Schultz, and S. Goss. 2008. Unfunded obligation and transition cost for the OASDI program. Actuarial Note no. 2008.1. Social Security Administration, Office of the Chief Actuary, July.