Volume Title: Social Security Policy in a Changing Environment

Volume Author/Editor: Jeffrey Brown, Jeffrey Liebman and David A. Wise, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 978-0-226-07648-5

Volume URL: http://www.nber.org/books/brow08-1

Conference Date: October 19-22, 2006

Publication Date: June 2009

Chapter Title: Pricing Personal Account Benefit Guarantees: A Simplified Approach

Chapter Author: Andrew Biggs, Clark Burdick, Kent Smetters

Chapter URL: http://www.nber.org/chapters/c4541

Chapter pages in book: (229 - 249)

7

Pricing Personal Account Benefit Guarantees
A Simplified Approach

Andrew Biggs, Clark Burdick, and Kent Smetters

7.1 Introduction

A number of proposals to introduce personal accounts to the Social Security program contain provisions that would guarantee account holders against relatively poor investment performance that would make their total benefits fall below the level scheduled under current law. Such protections are attractive to account participants, who would gain the financial and other potential advantages of personal accounts without the principal downside risk of relatively poor investment performance.[1] However, given the size of Social Security benefit entitlements and the potential risk of market investment, guarantees constitute a significant contingent liability to whomever would be providing the guarantee, whether it be the private markets or the government. For that reason, it is important to fully evaluate the potential costs of guaranteeing private investments against market risk.

Although some academic researchers, most notably George Pennacchi

Andrew Biggs is a resident scholar at the American Enterprise Institute. Clark Burdick is an economist in the Division of Economic Research, Office of Policy, at the U.S. Social Security Administration. Kent Smetters is an associate professor of insurance and risk management at The Wharton School, University of Pennsylvania, and a research associate of the National Bureau of Economic Research.

Smetters's research was supported by the U.S. Social Security Administration (SSA) through grant #10-P-98363-1-03 to the National Bureau of Economic Research (NBER) as part of the SSA Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA, any agency of the federal government, or of the NBER. The authors thank Steve Goss, Olivia Mitchell, David Pattison, and George Pennacchi for many helpful comments but maintain sole responsibility for any remaining errors or inaccuracies.

1. Examples of potential advantages are portfolio choice, inheritability, clear property rights, and so on.

(1999) and Mari-Eve Lachance and Olivia Mitchell (2003), have shown considerable interest in the market cost of benefit guarantees, most policy analysis has tended to focus on the expected or mean cost of personal account benefit guarantees. An expected cost approach evaluates the probabilities of various outcomes and reports back the average or "expected" cost of a guarantee provision.

Expected costs provide valuable information, but do not reflect the greater valuation placed by the market on losses relative to the expected value of the losses. Indeed, the Congressional Budget Office (2006a, 18) has recently shown that the total cost of a benefit guarantee, including the associated cost of market risk, could be as much as three times larger than its expected cost. A so-called risk-neutral valuation provides such market information and thus may be useful to policymakers. This chapter demonstrates how a model for calculating the expected cost of a benefit guarantee can easily be modified to present the market price of personal account guarantees as a supplement to expected cost valuations.

We begin with a discussion of proposals to incorporate personal retirement accounts (PRAs) into Social Security and why some proposals have included guarantees against adverse investment outcomes. We also discuss the current actuarial analysis of Social Security personal account guarantees, which reports the expected cost of such a guarantee.

We then outline a simple method for producing a market-priced cost estimation of a guarantee against relatively poor investment performance. It is first shown for a simple example of a stock purchase to illustrate that it produces results equivalent to the Black-Scholes model. We then outline how such an approach could be useful in evaluating guarantees for personal accounts, where using an explicit Black-Scholes approach can be cumbersome.

We illustrate our approach using a Social Security reform proposal from Senator John Sununu (R-NH) and Representative Paul Ryan (R-WI). This proposal would introduce personal accounts investing from 5 to 10 percent of wages, depending upon the worker's earnings level. At retirement, individuals would receive either the proceeds of their personal account or their currently scheduled benefit, whichever was greater. Thus, this plan effectively guarantees that accounts would produce benefits no lower than those scheduled for the current program.

We first construct a simple model to estimate the expected cost of the benefit guarantee in the Ryan-Sununu proposal. This model is calibrated to roughly replicate the expected cost estimates produced by the SSA Office of the Chief Actuary (OACT). We then make a simple alteration to this model to produce a risk-neutral estimate of the guarantee cost. Estimates of the market cost of the guarantee using a risk-neutral valuation derived from our preferred approach, a stochastic modeling exercise that uses carefully calibrated Monte Carlo simulations, are similar to the results of the simple model.

It is worth noting that in our simple model, a number of variables are not modeled stochastically, including wage growth and inflation. Hence, any correlation between career-length wage growth and market returns is precluded. We also exclude the possible effects of the presence of a guarantee on portfolio allocations over time. While these issues are important for the consideration of the costs of any guarantee, be it from the expected cost or market-price perspective, they do not weigh on the choice between these two perspectives.

We close with a discussion of outstanding issues regarding personal account guarantees.

7.2 Types of Benefit Guarantees

The Social Security program is projected to experience financial strains as the baby boom generation retires and the population ages. Social Security's Trustees project that the program cost will begin to exceed tax revenues in 2017 and that its trust fund will be exhausted in 2040. At that point, the program would be capable of paying around 74 percent of scheduled benefits, with larger reductions in future years (Social Security Administration 2006).

A number of proposals to reform Social Security for the future have incorporated PRAs, similar to simplified individual retirement accounts (IRAs) or 401(k) accounts. Under personal account plans, individuals would invest part of their existing payroll taxes, additional contributions, or tax credits funded by general revenues into accounts holding portfolios of stocks, corporate bonds, and government bonds. At retirement, the proceeds of the account would augment or replace benefits paid from the traditional program.

Personal accounts invested in equities will tend to increase average retirement benefits for workers choosing to participate. This is one of the principle reasons advocates favor personal accounts: the higher expected benefits they provide would generally cushion against reductions in traditional benefits that could be used to balance the program's finances. Critics charge that expected value analysis ignores risk. While people might do better by holding a personal account most of the time, they could actually do worse.

In response, over the years a number of Social Security proposals have contained guarantees against adverse market outcomes. A number of different types of guarantees are possible, including guarantees of minimum rates of return on account savings, guarantees against retiring in poverty, and so on. However, almost all actual reform proposals have guaranteed current law scheduled benefits.[2] That is, a worker with a personal account (or spouses or widows receiving auxiliary benefits based upon that worker's earnings) would be guaranteed at retirement a benefit at least as high as those scheduled under current law, or more if the account balance could provide it.[3]

7.3 Current Practice: Expected Costs

The expected cost of a Social Security reform proposal's guarantee against market risk can be estimated based upon assumptions regarding the expected rate of return and standard deviation of portfolios held in personal accounts. From these assumptions, the mean and distribution of account balances (and the annuities they can purchase) is estimated (or approximated) relative to the guaranteed level, which is generally current law scheduled benefits. From this distribution the percentage of accounts falling short of the guaranteed level is calculated, as well as the average amount by which such accounts fall short. The average shortfall across all outcomes (or across the entire distribution), which is, therefore, also the average payment to satisfy the guarantee, represents the expected average cost of the guarantee. Projected across the retiree population, an aggregate expected cost of the guarantee can be calculated.

7.3.1 Advantages of Expected Cost Analysis

The expected cost of a guarantee is useful for budgeting, which is a primary use of actuarial analysis of current law Social Security and alternative proposals. The expected cost constitutes a "best guess" of what a guarantee will cost in a particular year. The 1990 Federal Credit Reform Act requires that the future costs of certain guarantees, although not personal account guarantees explicitly, be recorded on the budget. The costs of these guarantees are also typically recorded on an expected value basis (with some exceptions) using discount rates and procedures provided by the Office of Management and Budget.

2. All current proposals containing guarantees would ensure that individuals receive at least the benefit scheduled under current law. However, it is also possible to provide rate of return guarantees for accounts. These could be relevant for proposals that "offset" traditional benefits based upon contributions to accounts compounded at a given rate of interest. A guarantee of that interest rate on account contributions would ensure that account holders receive no less in total benefits than had they not participated in an account. While scheduled benefit guarantees and rate of return guarantees differ in form, they are analytically similar. Simply put, a scheduled benefit guarantee merely guarantees that an account produce a return sufficient to purchase an annuity equal to the portion of scheduled benefits that would not be payable from the traditional program under the plan. This implicit return would be different for each individual, but there is no fundamental difference between the two approaches.

3. Even aside from protecting against market risk, this is a relatively generous guarantee for younger individuals given that under current law, benefits would be cut significantly from scheduled levels once the trust funds became exhausted (which is currently projected to occur in 2040).

7.3.2 Disadvantages of Expected Cost Analysis

For expected cost analysis to be a useful guide for policymakers, however, the underlying risk must be fairly diversifiable from the government's perspective. A diversifiable risk is both small and uncorrelated with the other risks in the economy, including the tax base. Under these conditions, the classic Arrow-Lind theorem (1970) showed that the government should essentially be indifferent to risk and, therefore, discount future risky liabilities by the risk-free rate. An analogous result appeared in the Capital Asset Pricing Model and related work around the same time (Borch 1962).

The government might be able to diversify some risks better than the private sector if private markets are "incomplete" in at least one of two ways. First, some households in the economy might be underexposed to market risks, perhaps due to various fixed costs associated with investing (Abel 2001) or myopia. Exposing these households to market risks, maybe with personal accounts containing little or no overhead costs, could potentially increase their welfare (Diamond and Geanakoplos 2003; Campbell et al. 1999).

However, a guaranteed benefit backed by the government undermines some of this risk sharing and instead transfers risk to workers by increasing their risk of tax increases. Unlike financial markets, though, the government cannot distinguish between workers who are willing to take this risk and those who are not. Some workers might be more tolerant to additional risk because of their preferences or if their human capital returns (wages) are minimally correlated with stock market returns. Spreading this risk indiscriminately throughout the entire economy could actually harm households. Indeed, given the low investment fees now being charged by the private sector, the government could presumably improve risk sharing using guarantees only if many households are myopic or uneducated about saving and investment.

Second, the government might be able to diversify risk across generations because it is impossible for the private sector to write risk-sharing contracts with the unborn (Bohn 2003). A benefit guarantee would naturally shift resources from older retirees to younger workers through the tax system. In essence, younger workers would get exposure to the stock market risk of the preceding generation, something that they could not do directly through capital markets. Connecting generations in this manner, therefore, could improve risk-sharing.[4]

4. Social Security benefit guarantees could be limited by trust fund solvency, such that if the program became insolvent the guarantee would not be honored. However, the Ryan-Sununu proposal examined here and some other Social Security proposals contain provisions for transfers of general tax revenue as needed to maintain solvency, implying that that guarantee would be honored even if the program required additional non-Social Security resources to do so.

However, this argument requires that the human capital returns (wages) of younger workers are not sufficiently correlated with the stock returns of the preceding generation. This assumption is difficult to test at a generational level because there are only three or four unique data points at such low frequency. Nonetheless, while the associated standard errors are large, the data seem to suggest that human and physical capital returns are highly correlated, at around 0.8. In the context of the neoclassical model, that means that low frequency shocks can mainly be traced to changes in productivity rather than depreciation. While a correlation of 0.8 might still leave some room for shifting capital risk from older retirees to younger workers, it could also mean that the optimal direction of risk-sharing is just the opposite: from younger workers to older retirees (Smetters 2003).

It is true that markets do not currently offer options of the duration necessary to guarantee lifetime accumulations in personal accounts. Moreover, the vast size of the guarantees necessary for Social Security guarantees make it unlikely the market alone could provide them. For this reason, it might be argued that pseudomarket prices are not relevant to guarantees that would most likely be offered by the government. However, this ignores the fact that government itself does not truly bear risks so much as spread them among the various parties who provide resources to or receive resources from the government. Hence, the market cost of risk remains relevant in evaluating the economic impact of a guarantee, even if private markets are not used to hedge the associated risk.

In summary, the consensus in the academic literature is that it is unlikely that the government has much, if any, advantage in risk-sharing relative to the private market. Moreover, even if the government did have an advantage, especially between generations, it is not obvious that the optimal direction of risk shifting is from older retirees to younger workers, as implicit in a benefit guarantee.

As a result, policymakers arguably should not treat risk much differently than individual investors who consider both expected outcomes and risk. For instance, an individual deciding his own 401(k) investment strategy would not focus solely on expected outcomes. Rather, he or she would consider that stocks, bonds, and other investments offer combinations of risk and return. Moreover, periods of low returns from a risky asset are likely to be correlated with poor outcomes in other areas, such as labor income.

7.4 Risk-Neutral Valuation

A risk-neutral valuation of a guarantee reflects the potential market cost for insuring against the underlying risk, which could augment expected cost analysis of this risk. For these reasons, academics and government agencies are increasingly calculating market valuations of contingent liabilities of the government.[5] Risk-neutral methods are a common approach for estimating the market price of transfers of risk.

To be sure, the expression "risk neutral" might be a bit confusing at first glance because it seems to indicate an indifference to risk. The rather arcane expression reflects the assumed efficiency of capital markets, that is, that they are complete. In this case, the private-market cost that might be charged for a benefit guarantee can be priced using no-arbitrage relationships with private-market assets, riskless transactions that do, however, reflect the premium over expected cost that is demanded by markets to cover the risk.

A simple example illustrates the differences between an expected cost and a risk-neutral cost valuation. Consider a provision in which personal accounts were invested in stocks and the government guaranteed account holders against any returns below the long-term average for stocks. In exchange, the government reclaimed or "clawed back" any returns above that long-term average. Assuming a normal distribution of returns, the reclaimed returns above the average should be sufficient to compensate account holders for returns below the average. Thus, the expected cost of such a guarantee is zero.

Financial markets, however, would charge a significant premium for such a guarantee because it guarantees the equity premium, that is, the difference between average equity returns and the risk-free rate paid by government debt. In fact, ignoring additional administration charges, the cost of this guarantee would be exactly equal to the equity premium itself. Investors, therefore, would be exactly indifferent between investing in government bonds versus investing in equities and purchasing a guarantee.

By the Hans Stoll put-call parity relationship, this type of guarantee can be decomposed into two transactions: give investors a put option that allows them to sell the stock at an exercise or "strike price" implied by its expected return, while requiring investors to sell a call option allowing the government the right to buy the stock at the same strike price. Because the strike price exceeds the price implied by appreciation of the stock at the *risk-free* rate, the underlying put option would be much more valuable than the call option. The two options would have equal value only if the strike price were tied to the risk-free rate, which is much less than the expected return to equities.

7.4.1 Black-Scholes

The Black-Scholes option pricing formula is probably the easiest way to compute the cost associated with put and call options using the no-arbitrage approach. The Black-Scholes price of a call option is equal to

5. Pennacchi (1999), Lachance and Mitchell (2003), and the Congressional Budget Office (2006a) apply risk-neutral methods to individual accounts. The Congressional Budget Office (2004, 2006b) applies risk-neutral methods to federal loan guarantees.

$$C_0 = S_0 N(d_1) - Xe^{-rt} N(d_2),$$

where

$$d_1 = \frac{\ln(S_0/X) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$$

and

$$d_2 = d_1 - \sigma\sqrt{T}$$

and

C_0 = the call option price
S_0 = the purchase price
$N(d)$ = the probability that a random draw from a standard normal distribution will be less than the value d
X = the exercise price
e = the base of the natural log function (2.71828)
r = the riskless rate of return
σ = the standard deviation of the log of gross portfolio returns
T = the length of the option, or the time until maturity

Then, the put-call parity relationship implies that the put option price is equal to

$$P = C_0 + \mathrm{PV}(X) - S_0 = C + Xe^{-rt} - S_0,$$

where P equals the put option price, and PV(X) equals the present value of the exercise price. It is worth noting that the expected return on the asset plays no role in the formula: the option price is derived solely from the volatility of the asset and the riskless return. A more detailed discussion can be found in Ingersoll (1987).

7.4.2 General Risk-Neutral Valuation

The Black-Scholes formula, though, does not easily accommodate investments that are made and accumulated in personal accounts over numerous working years. The Black-Scholes formula would treat contributions made to personal accounts in each year separately. A Social Security benefit guarantee, though, would be applied to the accumulation of assets over many years. For an investment in any given year, the Black-Scholes formula would not recognize the amount of past accumulations.

More general risk-neutral methods pioneered by Cox and Ross (1976), however, can easily accommodate this added complexity. Our approach follows Hull (2002):

1. Sample a random path in a risk-neutral world: generate a return path based upon the *risk-free* rate of return and the standard deviation of annual returns on the *risky* asset.

2. Calculate the payoff from the guarantee: if the end balance is below the guaranteed level, the payoff is positive.

3. Repeat steps 1 and 2 to get many samples of the payoff in a risk-neutral world.

4. Calculate the mean of the sample payoffs to get an estimate of the expected payoff in a risk-neutral world.

5. Discount the expected payroll at the risk-free rate to get an estimate of the value of the guarantee.

Multiple contribution dates can be easily incorporated in Steps 1 and 2.

In anticipation of where we are headed, expected cost analysis in essence already follows Steps 1 to 5 with one difference in Step 1: the random path is generated using a rate of return larger than the risk-free rate to incorporate some expected equity returns. Simply reducing this parameter to the risk-free rate would allow a correctly specified expected cost model to calculate the market value of the underlying risk.

Consideration of these steps reveals an additional advantage offered by a risk-neutral valuation. In addition to providing a market cost estimate, a risk-neutral valuation is generally considered less subjective and potentially more accurate than an expected cost approach.[6] Notice that an expected cost analysis requires knowledge or a forecast of expected equity returns in addition to all of the information required for a risk-neutral valuation. Hence, a risk-neutral valuation requires less uncertain, and potentially subjective, information. Merton (1980) discusses some of the difficulties in estimating the equity premium (or equivalently the expected return on equities).

7.4.3 Comparison of the Two Approaches

For a large number of simulated paths and a *single* contribution date, the Black-Scholes formula and the more general approach outlined in the preceding should produce the same value for a benefit guarantee.

To illustrate, consider an individual who purchases \$100 of stocks, with an expected return of 6.5 percent above inflation and a historical standard deviation of annual returns of 20.6 percentage points.[7] He or she intends to hold these stocks for ten years, with an expected end balance of \$187.71. ($\$100 \times 1.065^{10}$) However, he or she wishes to purchase a guarantee that he or she can sell his stocks for no less than that amount ten years hence.

Using the Black-Scholes formula, the cost of a put option guaranteeing that \$100 of stocks purchased today can be sold for \$187.71 in ten years

6. The authors are especially grateful to George Pennacchi for pointing this out.

7. Notice than an annual standard deviation of stock returns of 20.6 percent implies a sigma, or volatility, of 19.17 percent. That is the parameter sigma in the Black-Scholes option pricing formula refers to the standard deviation of the log of gross returns and not the annual standard deviation of returns.

time would be $51.94. This is an expensive guarantee, equal to over half the initial purchase price and 28 percent of the guaranteed end balance.

The alternate approach, outlined in the preceding, stochastically generates a number of outcomes, with the initial purchase price compounded at the riskless 3 percent rate of return and varying with the historical 20.6 percent standard deviation for stocks. Due to the lower assumed rate of return, the mean end balance after ten years of 500,000 simulations equals $134.47, with a standard deviation of $93.16. Of the end balances, 79 percent are below the guaranteed value of $187.71, with an average shortfall (including instances of no payout) of $71.47. The present value of this shortfall is $53.18, a difference of only 2 percent from the value derived with Black-Scholes. This difference is primarily due to sampling variation; increasing the number of sample paths would reduce the difference even more.

Repeating this exercise but accumulating balances using the expected annual return to equities of 6.5 percent rather than the risk free rate as the mean for the simulations generates much different results. The average balance in 500,000 simulations becomes $187.45 with a standard deviation of $124.73. The average shortfall across simulations is $44.74, which has a present value of $33.29. This $33.29 represents the expected cost of the guarantee under an expected cost approach. Clearly the risk-neutral cost of $53.18 is a much better estimate of the market cost of $51.94 implied by the Black-Scholes option pricing formula. The expected cost approach underestimates the market cost by 36 percent.

The advantages of the more general approach become more apparent when applied to Social Security personal accounts. In this case, the government is not providing a guarantee that a single purchase made on one date can be sold for a given price at a stated later date. Rather, individuals make a number of purchases throughout their lifetimes, on an annual or more frequent basis, the compounded sum of which must be sufficient to purchase an annuity equal to their scheduled Social Security benefits.

7.5 A Simple Risk-Neutral Valuation Model for Benefit Guarantees

This section develops a simple model to show how risk-neutral valuation can be used to estimate the market value of the underlying risk associated with a benefit guarantee. Our model first attempts to replicate the expected cost of Social Security guarantees as projected by the SSA's Office of the Chief Actuary. It then alters the parameters as outlined in the preceding to estimate a risk-neutral cost for an identical guarantee. This model is designed solely for illustrative purposes. The technique outlined in the preceding could easily be applied to more detailed microsimulation models, though with an increase in computation time.

To illustrate, we make these calculations for the Social Security reform proposal from Senator John Sununu (R-NH) and Representative Paul Ryan (R-WI). While several other reform plans include guarantees, the Ryan-Sununu proposal is relatively simple in its construction, making for ease of modeling, and has been scored by OACT, thereby providing a baseline to ensure that the simple model roughly replicates existing expected cost estimates (Goss 2005).

7.5.1 Ryan-Sununu Proposal

Once phased in, individuals under the Ryan-Sununu plan would have personal accounts investing 10 percent of taxable earnings up to $10,000 (indexed with wages from 2006) and 5 percent of taxable wages above that level. Accounts are assumed to be invested in a portfolio consisting of 65 percent stocks and 35 percent corporate bonds, with annual administrative costs equal to 0.25 percent of assets managed. Stocks are projected to earn 6.5 percent above inflation and corporate bonds 3.5 percent, for an expected return net of administrative costs of 5.2 percent above inflation.

At retirement, individuals would receive either the annuitized value of their PRA balance or their currently scheduled benefit, whichever was greater.[8] This guarantee effectively entails supplementing personal account balances that fall short of the level needed to purchase an annuity equal to current law scheduled benefits.

For the purposes of calibrating our model, we simulate individuals who spend a full working lifetime under the Ryan-Sununu proposal. This eliminates the need to model implementation provisions contained in the plan for individuals spending only part of their careers with personal accounts.

7.5.2 Outline of Our Model

We base our model on the stylized scaled earner patterns produced by SSA's Office of the Chief Actuary (Clingman and Nichols 2005). Scaled earners have a typical hump-shaped life-cycle pattern of earnings from age twenty-one through age sixty-four. These earnings patterns are derived from a longitudinal sample of historical earnings records and are commonly used to simulate the effects of changes to the benefit formula and the introduction of personal retirement accounts upon individuals. These earnings profile exhibit the typical inverted U-shaped pattern over the life cycle. A medium-scaled earner would begin his or her working career with earnings below the national average wage, have earnings above the national average in middle age, and then have declining relative earnings as he or she neared retirement. With the exception of the maximum wage

8. In fact, individuals would be required to purchase an annuity with their PRA balance providing benefits equal to those scheduled under current law. If the PRA balance exceeded the annuity cost, extra funds could be withdrawn as a lump sum. If the PRA balance was not sufficient to purchase the required annuity, the guarantee provision would supplement the PRA balance to the necessary level.

worker, scaled earners at higher or lower earnings levels follow the general pattern of the medium-scaled earner, though at different absolute levels of earnings. We consider five different scaled earnings patterns, plus a steady earner at the maximum taxable wage:

- Very low: lifetime earnings at the 13th percentile of the distribution.
- Low: lifetime earnings at the 27th percentile.
- Medium: lifetime earnings at the 57th percentile.
- High: lifetime earnings at the 82nd percentile.
- Maximum taxable wage: lifetime earnings at the 100th percentile.

Applying these scaling factors against the average wage index projected by the Social Security Trustees, we can produce simulated earnings and account contributions.

For each worker type, a projected personal account balance is calculated consistent with OACT methods, in which annual account contributions are compounded at the projected geometric mean return for the assumed account portfolio, minus administrative costs.[9] Expected account balances at age sixty-five are converted to annuities based upon mortality and interest rate projections from the Social Security trustees.

The key statistic for this model's distribution of account balances is the coefficient of variation of final account balances, that is, the standard deviation of account balances divided by the mean balance. In lieu of a stochastic simulation, the variation in total account balances in retirement is estimated as the summed variation of account investments made in each year. Using a medium-scaled earnings pattern, account contributions for each year are calculated and individually compounded to age sixty-five at the mean expected return for the portfolio. The sum of these compounded annual contributions equals the expected account balance at retirement. Based upon the standard deviation of annual returns, the standard deviation of returns from the year a contribution is made through retirement is calculated for each year's contribution. An end balance is calculated for each year's contribution at the mean return minus the standard deviation of holding period returns; the sum of these balances represents the account balance at 1 standard deviation below the mean. The difference between this end balance and the end balance calculated at the expected return is the standard deviation; relative to the expected end balance, this difference is the coefficient of variation.

We calculate these values for the Ryan-Sununu default portfolio of 65 percent equities and 35 percent corporate bonds. The geometric mean returns are 6.5 percent and 3.5 percent, respectively, based upon standard Office of the Actuary projections. The standard deviations of returns and

9. The OACT also projects guarantee costs on a basis of an all-bond portfolio, as well as occasionally based upon a higher-yield assumption.

covariances between returns are from the 2006 Ibbotson yearbook (Ibbotson Associates 2006). The standard deviation of annual stock and corporate bond returns is taken to be 20.2 percent and 8.5 percent, respectively, and the correlation between them 0.19. Based upon the preceding method and these assumptions, the coefficient of variation for a personal account holding the Ryan-Sununu portfolio would be 50 percent.

Scheduled benefits at age sixty-five are calculated for each worker type. However, the Ryan-Sununu proposal guarantees all scheduled benefits, including auxiliary benefits paid to spouses and other eligible family members. The SSA Modeling Income in the Near Term (MINT) model projects that in 2050, auxiliary benefits will make up roughly 5 percent of total benefits paid to individuals of retirement age.[10] For that reason, scheduled benefits in 2050 are adjusted upward by 5 percent in an attempt to account for this provision.

Based upon scheduled benefits, expected account balances, and the distribution of account balances, we calculate the percentage of accounts for each worker type that could be expected to fall short of scheduled benefits and the size of the typical guarantee payment needing to be made.[11]

The next step is to convert benefits and guarantee estimates for each of these stylized workers into an approximation of costs covering the full population. This is accomplished using figures from OACT showing the percentage of individuals in the population who are best represented by each stylized worker type.

7.5.3 Calculation of the Expected Cost of Guarantee

Table 7.1 reports that 20.7 percent of the retiree population has average indexed monthly earnings (AIME) closest to those of the stylized very low earner; 22.4 percent closest to the low earner; 27.1 percent closest to the medium earner; 20.8 percent closest to the high earner; and 8.9 percent closest to the maximum wage earner.

The expected guarantee payment and scheduled benefit for each worker type are multiplied by the weighting factor. The sum of weighted guaran-

10. Calculations by SSA Office of Policy staff.

11. Note that correlation between market returns and wage growth could reduce personal account guarantee costs by a more mechanical route. Under current law, initial Social Security benefits are indexed to the growth of wages. If lifetime wage growth and market returns tend to be correlated, then individuals with low market returns would also tend to have low scheduled benefits, thereby reducing the cost of a personal account guarantee. Preliminary calculations (not shown here) by one author indicate that if working lifetime wage growth and market returns are perfectly correlated, the expected cost of a personal account benefit guarantee would decline by roughly one-quarter versus if lifetime wage growth and market returns are uncorrelated. If the correlation were 0.5, guarantee costs would decline by roughly one-eighth. Note, however, that this issue does not touch on the question of whether the expected cost or risk neutral valuation best expresses the value of the contingent liability to the guarantor. Rather, if correlation between wage growth and market returns is assumed, either an expected cost or a risk-neutral model should account for it.

Table 7.1 **Calculation of expected guarantee costs for individuals retiring at age 65 in 2050**

Earnings level	Very low	Low	Medium	High	Maximum
Average wage (wage indexed to present; US$)	8,516	15,329	34,065	54,112	72,342
Percentile of earnings distribution	13.4	27.1	57.4	82.1	100.0
Percent of workers closest to stylized worker	20.7	22.4	27.1	20.8	8.9
Scheduled benefits (US$)	9,808	12,832	21,138	28,024	34,568
Adjusted scheduled benefits (US$)	10,347	13,538	22,301	29,565	36,469
Expected annuity from personal account (US$)	9,969	15,031	31,545	38,273	64,245
Standard deviation of personal retirement account annuities (US$)	4,985	7,516	15,773	19,137	32,123
Percentage of account holders accessing guarantee	53	42	28	32	19
Average guarantee payment (US$)	1,859	1,917	2,185	3,292	2,707
Average guarantee payment as percentage of average benefits	18	14	10	11	7
Weighted value of benefits (US$)	2,142	3,032	6,043	6,150	3,246
Weighted value of average guarantee payment (US$)	385	429	592	685	241
Guarantee cost as percentage of total benefits	11.3				

Source: Authors' calculations.

tee payments is then expressed as a percentage of the sum of weighted benefit payments. Under these calculations, expected guarantee payments would equal roughly 11.3 percent of total benefits to new retirees in 2050.

According to the OACT analysis of the Ryan-Sununu proposal, expected guarantee costs in 2050 would equal $190.3 billion (in $2004). This amount is equal to 13.3 percent of total OASI costs in 2050, based upon projections from the 2004 Social Security trustees report. Thus, our simple model's estimates appear sufficiently close to proceed to the next step of converting the expected cost of the guarantee to the risk-neutral cost.

To repeat, this model is not intended to replicate existing results with precision, particularly as parameters used may be different. Nevertheless, as a calibrated replication of current results, it gets close enough to projected costs to illustrate the effects of the modified parameter input we propose here.

7.5.4 Calculation of the Market Cost of the Underlying Risk

Now that we confirmed that our model produces an estimate of the expected costs of a benefit guarantee that is roughly consistent with existing estimates, we then alter the model in order to estimate guarantee costs on a risk-neutral basis. The single change to the model's inputs is that the mean account balance is now produced by compounding account contributions at the rate of return projected to be earned by the Social Security trust funds rather than the expected return from the stock-corporate bond portfolio used in estimating expected guarantee costs. However, as detailed in the preceding the distribution of account balances expressed through the coefficient of variation remains the same as with the risky portfolio.[12]

The merit of this approach is that this conversion consists solely of altering the distribution of account balances at retirement from one based upon the expected return to one based upon the riskless return. That is, expected PRA balances at retirement are lower, but all other parameters remain the same. Thus, as detailed in table 7.2, projected end balances compounded at the bond rate are considerably lower than when compounded at the expected return, equaling roughly 60 percent of the expected account balance.

As expected, compounding returns at a lower rate of return increases the proportion of account holders whose balances require access to the guarantee and the size of the average guarantee payment. The guarantee cost relative to total benefits rises from 11.3 percent under expected cost valuation to 28.2 percent under risk neutral valuation, a factor of 2.5. If these proportions held true throughout the seventy-five-year scoring period, the present value expected guarantee cost of slightly over $2 trillion would rise to almost $5 trillion.

7.5.5 Change in Portfolio Composition

In proposals that allow for portfolio choice, it could be expected that inclusion of a benefit guarantee would alter the average portfolio allocation of account holders. In essence, the account holder is given two things of value: a cash allotment to be invested in the account and an implicit put option against losses relative to a given baseline. The present value of the account contribution is the same regardless of what it is invested in. The value of the option, however, rises with the volatility of the chosen portfolio. For that reason, rational account holders would tend to increase the share of

12. Note that the choice as the riskless rate of the projected return on the Social Security trust funds, whose special issue assets earn interest rates equal to the average of medium- and long-term government bonds in the market, will produce lower projected guarantee costs than a short-term bond rate, which might be more accurately described as riskless.

Table 7.2 **Calculation of expected guarantee costs for individuals retiring at age 65 in 2050 (risk-free return)**

Earnings level	Very low	Low	Medium	High	Maximum
Average wage (wage indexed to present; US$)	8,516	15,329	34,065	54,112	72,342
Percentile of earnings distribution	13.4	27.1	57.4	82.1	100.0
Percent of worke0rs closest to stylized worker	20.7	22.4	27.1	20.8	8.9
Scheduled benefits (US$)	9,808	12,832	21,138	28,024	34,568
Adjusted scheduled benefits (US$)	10,347	13,538	22,301	29,565	36,469
Expected annuity from personal account (US$)	6,071	9,072	15,879	23,181	37,768
Standard deviation of personal retirement account annuities (US$)	3,036	4,536	7,940	11,591	18,884
Percentage of account holders accessing guarantee	92	84	79	71	47
Average guarantee payment (US$)	4,215	4,519	6,744	7,593	5,796
Average guarantee payment as percentage of average benefits	41	33	30	26	16
Weighted value of benefits (US$)	2,142	3,032	6,043	6,150	3,246
Weighted value of average guarantee payment (US$)	873	1,012	1,828	1,579	516
Guarantee cost as percentage of total benefits	28.2				

Source: Authors' calculations.

stocks in their account portfolios in response to a benefit guarantee. The option value of the account would be maximized if invested solely in stocks.

While limits on portfolio allocations could be implemented to control for such effects in personal account plans, the Ryan-Sununu proposal does not limit the guarantee contingent upon holding a specified portfolio. Initially three portfolios would be offered, with equity components of 50, 65, and 80 percent, respectively, with the remainder held in corporate bonds. Once account balances reached $2,500 (in 2005 dollars), additional investment options would be available through private investment companies.

To illustrate the potential cost effects of changing portfolio allocations on guarantee costs, we repeat our preceding calculations for the Ryan-Sununu proposal but assume a portfolio of 100 percent stocks. Changing to an all equity portfolio does not alter any parameters other than the coefficient of variation of PRA annuities, which rises from 50 percent under the default 65-35 portfolio to 61 percent with all stocks. Doing so increases the risk-neutral cost of the Ryan-Sununu guarantee from 32.4 percent of total benefits to 33.9 percent. This relatively modest increase is due to the fact that the default portfolio already contains 65 percent equities, so the variance of outcomes does not increase a great deal.

However, larger costs are possible if account holders choose to vary their portfolios to "time the market." This could potentially increase costs further depending upon how this timing affected the variability of account portfolios. One advantage of the approach we introduce here is that if the effects of variable portfolio allocations are modeled for the purposes of calculating the expected cost of a personal account guarantee, those effects would be similarly treated in calculating the market cost of such a guarantee. That is, the change to parameter inputs we introduce to convert expected cost projections to risk-neutral valuations is not contingent upon modeling variability of account portfolios.

7.5.6 Alternate Calculations

As a check on the simple model presented in the preceding, we recalculate risk-neutral guarantee costs using a stochastic model, which is a preferred methodology for such an exercise. Nominal earnings profiles are created beginning in 2006 using the AWI and standard scaled earner profiles for very low, low, medium, and high earners. An additional nominal earnings profile is created for maximum earners who earn the taxable maximum in each year. We create earnings profiles for each of the thirty-one age cohorts who work a full forty-four years (ages twenty-one to sixty-four) between 2006 and 2079. Each of these cohorts is assumed to retire at age sixty-five in the years 2050 to 2080.

For workers at each age, nominal account contributions are calculated consistent with the Ryan-Sununu specifications. Nominal contributions are converted to constant 2004 dollars using the Consumer Price Index (CPI) assumptions from the 2004 Old-Age, Survivors, and Disability Insurance (OASDI) trustees' report. The real contributions are then accumulated at stochastic real annual rates of return less administrative costs equal to 25 basis points (0.25 percent). The stochastic real gross rates of return are generated from independently and identically distributed (i.i.d.) lognormal random variables to form 10,000 rate of return paths, each path representing the seventy-four years from 2006 through 2079. The lognormal rates of return are calibrated to have an expected value in levels equal to the OACT assumptions for the annual return on a PRA portfolio invested 65 percent in stocks and 35 percent in corporate bonds.[13] The standard deviation of the stochastic rates is assumed to equal 12.59 percent, the historical standard deviation of annual returns for a portfolio invested 65 percent in the S&P 500 with dividend reinvestment and 35 percent in a AAA corporate bond index.[14]

The accumulation of real (constant 2004 dollars) contributions at real annual rates of return results in a distribution of real final PRA balances for each of the 31 cohorts in the year they turn 65. The distribution of PRA balances for each cohort is then compared with the cost of purchasing an inflation indexed annuity that pays the age sixty-five retirement benefit in that and all subsequent years. The guarantee is a one-time "top up" payment made to individuals in the year of their retirement whenever their final PRA balance is insufficient to purchase the current law benefit annuity.

For each type of worker, the expected cost of the guarantee is approximated simply as the arithmetic mean of the guarantee payments across the 10,000 stochastic simulations. An aggregate expected cost estimate is produced by expressing the guarantee cost as a percentage of current law scheduled benefit payments (the cost of the inflation indexed annuity) and weighting this cost for each type of worker by the population percentage most closely resembling that type of worker.[15]

Given this procedure for producing expected cost estimates for a benefit guarantee, obtaining a market-cost estimate, or risk-neutral valuation, is remarkably easy. To produce a market-cost estimate, we follow exactly the same procedure described in the preceding, except that the stochastic rates of return are calibrated to have an expected value equal to the real new issue rate of 2.9 percent for government bonds issued to the OASDI trust funds as assumed in the 2004 trustees' report.[16] Importantly, the variance of the stochastic rates of return is unaltered from the value used in producing the expected cost estimate in the preceding. That is, the expected return of the PRA portfolio is altered, but the assumed variability of the portfolio is not.

Everything else proceeds exactly as before. Real PRA contributions are

13. We also calibrated the stochastic rates to have a geometric mean equal to the Commission to Strengthen Social Security (CSSS) assumptions, but those results are not reported here.

14. Historical data for the S&P 500 and the AAA-rated corporate bond index from 1914 through 2005 were obtained from the Total Return Database of Global Financial Data Inc.

15. Note that this relies on a cohort measure of benefit cost and not on a calendar year measure as is generally reported in the OASDI trustees' report. For a fully phased-in system of guaranteed personal accounts, there should be little difference between the guarantee cost as a percentage of cohort benefits and as a percentage of calendar-year benefit payments.

16. In 2004, the OASDI trustees assumed a nominal new issue rate of 5.8 percent and inflation of 2.8 percent annually under alternative II implying a real new issue rate of 2.92 percent. We use the new issue rate for trust fund assets assumed in the 2004 OASDI trustees' report as a proxy for the risk-neutral rate of return.

Table 7.3 **Risk neutral average guarantee payment as percent of average scheduled benefit (2050)**

	Very low	Low	Medium	High	Maximum
Stochastic model	41	34	32	27	17
Analytic model	41	33	30	26	16

accumulated at the stochastic risk-neutral rates, and the resulting distribution of real PRA balances is compared with the same inflation indexed annuity cost as before. Guarantee payments are again determined, but unlike before, the arithmetic average of these payments across all 10,000 simulations is now an approximate market cost for the guarantee. That is, the average guarantee payment using the stochastic risk-neutral rates approximates what an individual would pay in a competitive insurance market to purchase the benefit guarantee that ensures a minimum annual benefit equal to current law scheduled benefits.

Table 7.3 presents results from the stochastic model compared to those from the analytic model outlined in the preceding. In each case, the guarantee cost is represented as a percentage of average benefits for each worker type. These percentages are weighed to approximate costs for the full population.

As reported in the preceding, the analytic model calculates the risk-neutral guarantee cost for the Ryan-Sununu proposal in the year 2050 as 28.2 percent of total OASI benefits in that year. Using the same general method as the analytic model but inputs from the stochastic model generates an estimated risk neutral guarantee cost in 2050 of 29.6 percent of OASI benefits. Note, however, that the stochastic model does not contain an adjustment for auxiliary benefits. When that adjustment is dropped from the analytic model, the risk-neutral cost then declines to 25.8 percent of total OASI benefits. While neither model is a substitute for a full simulation against a representative population, they produce results roughly consistent with each other.

7.6 Conclusions

Once an appropriate model is constructed to calculate the expected cost of a guarantee, a change of a single parameter of that model enables the analyst to calculate the risk neutral guarantee cost as well. Our preferred approach uses a stochastic model to estimate the market value of the guarantee, but the risk-neutral price based on the analytic perspective turns out to be similar for the proposal modeled. From a practical perspective, the risk-neutral guarantee costs allows for greater information to be provided to policymakers with relatively little additional research cost.

References

Abel, Andrew B. 2001. The effects of investing Social Security funds in the stock market when fixed costs prevent some households from holding stocks. *American Economic Review* 91 (1): 128–48.

Arrow, Kenneth J., and Robert C. Lind. 1970. Uncertainty and the evaluation of public investment decisions. *American Economic Review* 60:364–78.

Bohn, Henning. 2003. Intergenerational risk sharing and fiscal policy. University of California, Santa Barbara. Mimeograph.

Borch, Karl H. 1962. Equilibrium in a reinsurance market. *Econometrica* 30: 424–44.

Campbell, John Y., Joao F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout. 1999. Investing retirement wealth: A life-cycle model. NBER Working Paper no. 7029. Cambridge, MA: National Bureau of Economic Research.

Clingman, Michael, and Orlos Nichols. 2005. Scaled factors for hypothetical earnings examples under the 2005 trustees report assumptions. Actuarial note no. 2005.3. Office of the Chief Actuary, Social Security Administration, September. http://www.ssa.gov/OACT/NOTES/ran3/an2005-3.html.

Congressional Budget Office. 2004. Estimating the value of subsidies for federal loans and loan guarantees. CBO Study. Washington, DC: CBO, August.

———. 2006a. Evaluating benefit guarantees in Social Security. Background Paper. Washington, DC: CBO, March.

———. 2006b. Valuing the student loan consolidation option. Background Paper, Washington, DC: CBO, October.

Cox, John E., and Stephen A. Ross, 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3:145–66.

Diamond, Peter, and John Geanakoplos. Social Security investment in equities. *American Economic Review* 93 (4): 1047–74.

Goss, Stephen C. 2005. Estimated financial effects of the Social Security Personal Savings Guarantee and Prosperity Act of 2005. Office of the Chief Actuary, Social Security Administration, April. http://www.ssa.gov/OATC/Solvency/RyanSununu_20050420.html.

Hull, John C. 2002. *Options, futures, and other derivatives.* 5th ed. Upper Saddle River, NJ: Prentice Hall.

Ibbotson Associates. 2006. *Stocks, bonds, bills, and inflation: 2006 yearbook.* Chicago: Ibbotson Associates.

Ingersoll, Jonathan. 1987. Option pricing theory. *The New Palgrave: A Dictionary of Economic Theory and Doctrine.* New York: Macmillan.

Lachance, Marie-Eve, and Olivia Mitchell. 2003. Guaranteeing individual accounts. *American Economic Review* 93 (2): 257–60.

Merton, Robert C. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8:323–61.

Pennacchi, George G. 1999. The value of guarantees on pension fund returns. *Journal of Risk and Insurance* 66 (2): 219–37.

Smetters, Kent. 2003. Trading with the unborn. NBER Working Paper no. 9412. Cambridge, MA: National Bureau of Economic Research.

Social Security Administration. 2006. *2006 annual report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance trust funds.* Washington, DC: GPO.

Comment George G. Pennacchi

This chapter by Andrew Biggs, Clark Burdick, and Kent Smetters makes two simple but very important points. First, if one has a model that can compute the expected cost of a personal retirement account (PRA) guarantee, then with a couple of changes in parameter values, the model can also compute the market cost of the guarantee. Second, knowledge of the guarantee's market cost is critical for determining sensible policy.

I agree wholeheartedly with these two results. In these comments, I will offer more intuition for the chapter's findings and add arguments for why policy should be guided by market costs and not expected costs. I will close with some suggestions for improving estimates of the market cost of PRA guarantees.

Biggs, Burdick, and Smetters construct a simple model that is calibrated to roughly replicate the Social Security Administration Office of the Chief Actuary's (OACT) expected cost of the Ryan-Sununu PRA guarantee. For a PRA invested 65 percent in stocks and 35 percent in bonds and assuming the expected real returns on stocks and bonds are 6.5 percent and 3.5 percent, respectively, they calculate that the guarantee's expected cost in 2050 equals 11.3 percent of total Social Security Old-Age and Survivors Insurance (OASI) costs. The expected cost in 2050 of the same guarantee computed by OACT's model is 13.3 percent of total OASI costs.

Having shown that their model is comparable to that of the OACT, they consider what would be the market cost of the same Ryan-Sununu guarantee, rather than its expected cost. Their computation of market cost is based on standard asset pricing methodology that accounts for the systematic (priced) risks inherent in stock and bond returns. Specifically, they take the identical model that was used to compute the expected cost of a PRA guarantee and alter two parameter inputs. Rather than setting the expected real returns on stocks and bonds equal to their physical (actual) values of 6.5 percent and 3.5 percent, respectively, they set them equal to their risk-neutral values, which for both is assumed to equal a risk-free real

George G. Pennacchi is a professor of finance at the University of Illinois.

This research was supported by the U.S. Social Security Administration (SSA) through grant #10-P-98363-1-03 to the National Bureau of Economic Research (NBER) as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the author and do not represent the views of the SSA, any agency of the federal government, or of the NBER.