

Pricing Personal Account Benefit Guarantees: A Simplified Approach

by

Andrew Biggs, Clark Burdick, and Kent Smetters

Discussion

by

George G. Pennacchi

**Professor of Finance
University of Illinois
1206 S. Sixth Street
Champaign, Illinois 61820
Tel.: (217) 244-0952
Email: gpennacc@uiuc.edu**

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

This paper 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 paper'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's) expected cost of the Ryan-Sununu PRA guarantee. For a PRA invested 65 % in stocks and 35 % in bonds and assuming the expected real returns on stocks and bonds are 6.5 % and 3.5 %, respectively, they calculate that the guarantee's expected cost in 2050 equals 11.3 % of total Social Security (OASI) costs. The expected cost in 2050 of the same guarantee computed by OACT's model is 13.3 % 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 % and 3.5 %, respectively, they set them equal to their risk-neutral

values, which for both is assumed to equal a risk-free real return of 2.9 %.¹ With this simple change, they find that the PRA guarantee's market cost in 2050 equals 28.2 % of total OASI costs.

Intuitively, it is not surprising that the market cost of this long-dated guarantee is almost two and one-half times its expected cost. The guarantor bears undiversifiable or systematic risk, because equity losses (and to a lesser extent bond losses) over a long horizon would tend to occur following several years of economic recession. This is precisely the scenario when the guarantor would least prefer to fund claims on PRA losses. A market cost for the guarantee that exceeds its expected cost reflects the market compensation for this systematic risk.

Because the systematic risk of a PRA guarantee derives from the systematic risk premia of the stocks and bonds in the PRA portfolio, a mathematically equivalent way of adjusting for this risk is to simulate scenarios of stock and bond returns using their risk-neutral expected returns, not their physical expected returns. Specifically, the correct market cost is found by simulating all security returns assuming that their expected return equals that of a short-maturity, risk-free, real return, such as the yield on a short-maturity Treasury Inflation Protected Security (TIPS). This lowers the simulated returns of the PRA portfolio because they grow, on average, at the 2.9 % risk-free rate rather than the portfolio's 5.45 % physical rate. Hence, more of these randomly simulated risk-neutral PRA returns result in a claim on the guarantor, thereby increasing the guarantee's cost.

The preceding description points to an advantage of valuing guarantees based on their market, rather than expected, costs. Estimates of market costs involve less subjectivity compared to estimates of expected costs. Computing market costs requires knowledge of the real return on a short-maturity, inflation-indexed bond. This real return is practically observable and can be

¹ This implies a risk-neutral expected PRA return (before administrative costs) of 2.9 % compared to a physical expected return of $(0.65 \times 6.5\% + 0.35 \times 3.5\%) = 5.45\%$. The paper's 2.9 % return is based on the real return that nominal Government bonds issued to the OASDI trust funds are expected to earn. Theoretically, this real return should be set to the real yield on a short-maturity, inflation-indexed Government bond. Doing so may lead to a risk-neutral real return somewhat lower than 2.9 % and result in a 2050 market cost slightly higher than what the paper calculates.

accurately measured from TIPS yields. In contrast, calculating expected costs requires knowledge of the physical expected returns on each of the securities in the PRA portfolio. Determining these expected returns over the life of the guarantee is more subjective, largely because it is not clear how they should be estimated. As discussed in Merton (1980), it is quite difficult to accurately estimate the expected returns on equities from their realized (historical) returns.² Such an exercise becomes even more difficult if one (realistically) assumes that expected equity returns vary over time.

Survey evidence suggests that real expected equity returns are time varying and that the equity risk premium, defined as the difference between the expected return on equity and the short-maturity, risk-free interest rate, has been falling. Figure 1 gives the median forecasts of the real returns on stocks, bonds, and bills, as well as the equity risk premium, from the Survey of Professional Forecasters (SPF) for the years 1992 to 2006.³ Except for the real return on Treasury bills, these median forecasts have tended to decline over this fifteen-year period. In particular, the median forecast for the equity premium was 5.0 % in 1992, peaked at 5.85 % in 1994, and is now down to 2.75 %. Further, the current median forecast for the real return on stocks over the next decade is 4.5 %, which is 200 basis points lower than the OACT's assumed real return of 6.5 %. An implication is that the OACT's estimate for the expected cost of a PRA guarantee is understated.

The SPF can be used to illustrate the subjectivity in determining the equity risk premium (and equity expected returns). Figure 2 plots each of the individual forecasters' estimates of the 10-year horizon equity premium. It is clear that there is wide cross-sectional variation in what

² Equity variances and covariance can be estimated relatively more accurately. The accuracy of variance-covariance estimates improve with the frequency at which returns are sampled, whereas the accuracy of expected return estimates increase with the return sample's observation period. It may take an observation period of 50 years or more to obtain an accurate estimate of a stock's expected return, assuming its expected return is constant over time.

³ The Federal Reserve Bank of Philadelphia has conducted the Survey of Professional Forecasters since 1990. It was previously carried out by the American Statistical Association and the National Bureau of Economic Research beginning in 1968. In recent years, approximately 30 professional economic forecasters have participated in this survey. See Croushore (1993) for details.

these professional forecasters believe will be the average excess return on stocks over the next decade. This suggests much potential disagreement among reasonable individuals as to the appropriate expected return estimates to use in calculating the expected cost of a PRA guarantee. Consequently, the potential for political manipulation of PRA guarantees costs is greater when they are based on expected costs rather than market costs that do not require knowledge excess expected returns.

Another reason for valuing PRA guarantees using market, rather than expected, costs is that policy based on expected costs can lead to large economic distortions. All else equal, expected costs of PRA guarantees are lower the greater is the systematic risk of the PRA portfolio. Specifically, consider two PRA portfolios having the same real return variance, but different systematic risk. The market costs of a PRA guarantee would be the same for these two portfolios, but their expected costs would differ. The portfolio with greater systematic risk, and therefore a greater systematic risk premium (excess expected return), would have a lower expected cost. Thus, policy based on expected costs produces an incentive for choosing PRA portfolios that are highly pro-cyclical. An implication is that the Government would face large guarantee claims following several years of economic recession, a time when Government budget deficits would already be large. Raising taxes or further increasing the Government's debt during such a scenario would be especially painful to current and future taxpayers. This incentive to amplify the magnitude of business cycles is avoided by use of market costs.⁴

Related to how a PRA guarantee should be valued is the issue of how it should be funded. Suppose that a premium that covered the cost of a PRA's guarantee was deducted from the account balance. Then based on the previous logic, setting premiums equal to the expected cost of the guarantee would induce individuals to choose a portfolio with excessive systematic risk. This moral hazard incentive is avoided by setting the premium equal to the guarantee's market

⁴ Other government guarantee programs that charge premiums equal to expected costs may also be subject to moral hazard that results in excessive systematic risk. Pennacchi (2006) describes how this distortion arises when bank deposit insurance premiums are based on expected, rather than market, costs.

cost. However, while a policy of market cost premiums avoids distortions, it may be politically problematic. Charging a premium above the guarantee's expected cost means that the Government guarantor would profit, on average, even though this profit reflects fair compensation to taxpayers for their exposure to systematic risk. To the financially unsophisticated, it might appear that the Government is overcharging for PRA guarantees.

A potential solution is to avoid explicit premiums but constrain PRA investment allocations along the lines proposed in Feldstein (2005). This involves investing a portion of the PRA portfolio in assets that pay risk-free real returns (TIPS) until the portfolio's guaranteed return is obtained. Under this portfolio restriction, the PRA guarantee is implicitly funded by the PRA holder who relinquishes the potential up-side returns that would be available from a riskier PRA portfolio.

I conclude by suggesting some extensions to the paper's simulation model that might fine-tune the estimates of the market cost of a PRA guarantee.⁵ In addition to stochastic security returns, I believe a model should recognize at least two other sources of uncertainty affecting PRA guarantees. One is the stochastic nature of wage contributions to a worker's PRA. A satisfactory simulation model should include risk-neutral processes for workers' real wages, where these wage processes might differ depending on a worker's age and industry.

Another potentially important source of uncertainty is the term structure of real interest rates. Figure 3 shows that TIPS yields have changed considerably over the past few years. Randomness in real interest rates affects guarantee costs in at least two ways. First, they cause variation in the risk-neutral expected returns on PRA securities, adding more randomness to a worker's PRA balance at retirement. Second, the real interest rate that a worker faces at retirement affects the level of annuity payments that can be purchased with her balance. If a PRA guarantee is truly a guarantee of minimum annual retirement benefits, rather than simply a final

⁵ These extensions are detailed in Pennacchi (1999) where they are incorporated into a model for valuing a Chilean PRA guarantee.

account balance guarantee, then the guarantee's cost depends on the term structure of real interest rates at the worker's retirement date.

A risk-neutral process for real interest rates based on the Vasicek (1977) model can be combined with risk-neutral processes for security returns and real wages. This can be done while assuming a general correlation structure between these processes. PRA guarantee costs are computed using Monte Carlo simulations of these processes for different worker cohorts.

References

- Croushore, Dean, 1993 "Introducing: The Survey of Professional Forecasters," *Business Review*, *Federal Reserve Bank of Philadelphia* (November/December), 3-13.
- Feldstein, Martin, 2005 "Reducing the Risk of Investment-Based Social Security Reform," National Bureau of Economic Research Working Paper 11084.
- Merton, Robert, 1980 "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics* 8, 323-61.
- Pennacchi, George, 1999 "The Value of Guarantees on Pension Fund Returns," *Journal of Risk and Insurance* 66, 219-237.
- Pennacchi, George, 2006 "Deposit Insurance, Bank Regulation, and Financial System Risks," *Journal of Monetary Economics* 53, 1-30.
- Vasicek, Oldrich, 1977 "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics* 5, 177-88.

Figure 1

Survey of Professional Forecasters Median 10-Year Forecasts

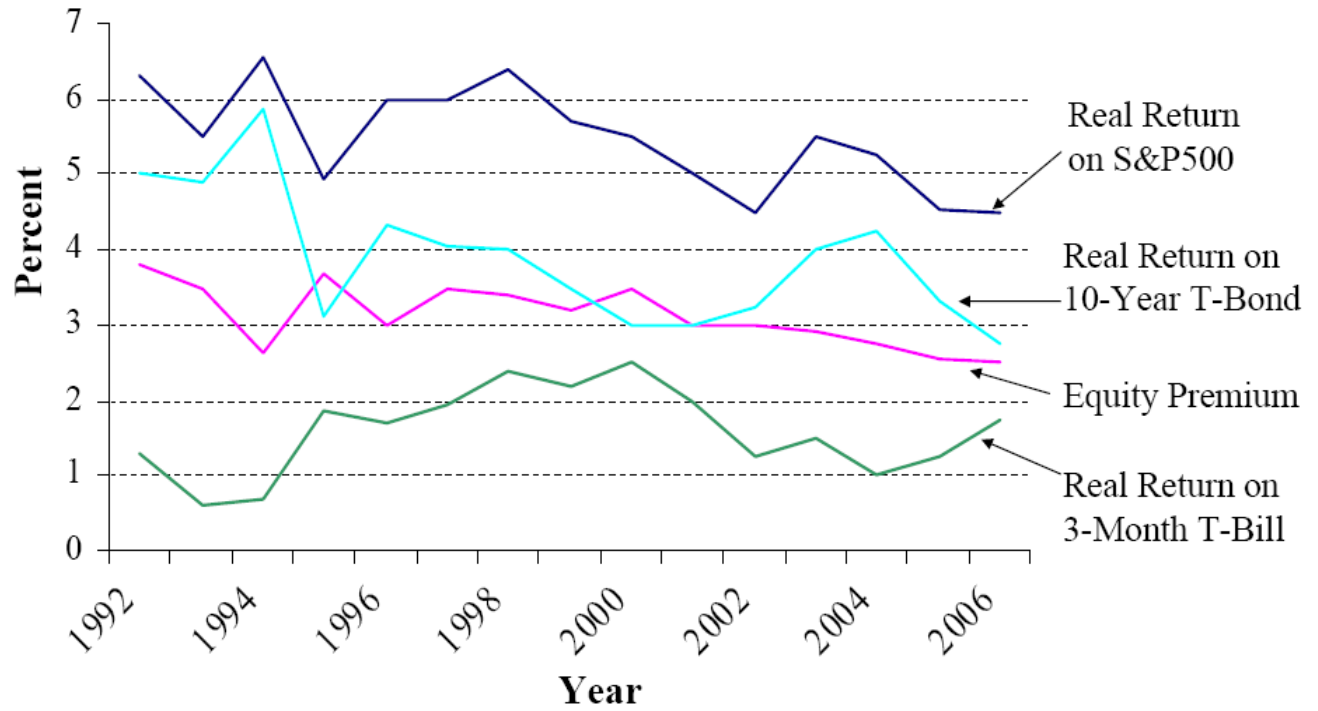


Figure 2

Survey of Professional Forecasters' Individual Equity Premium Forecasts

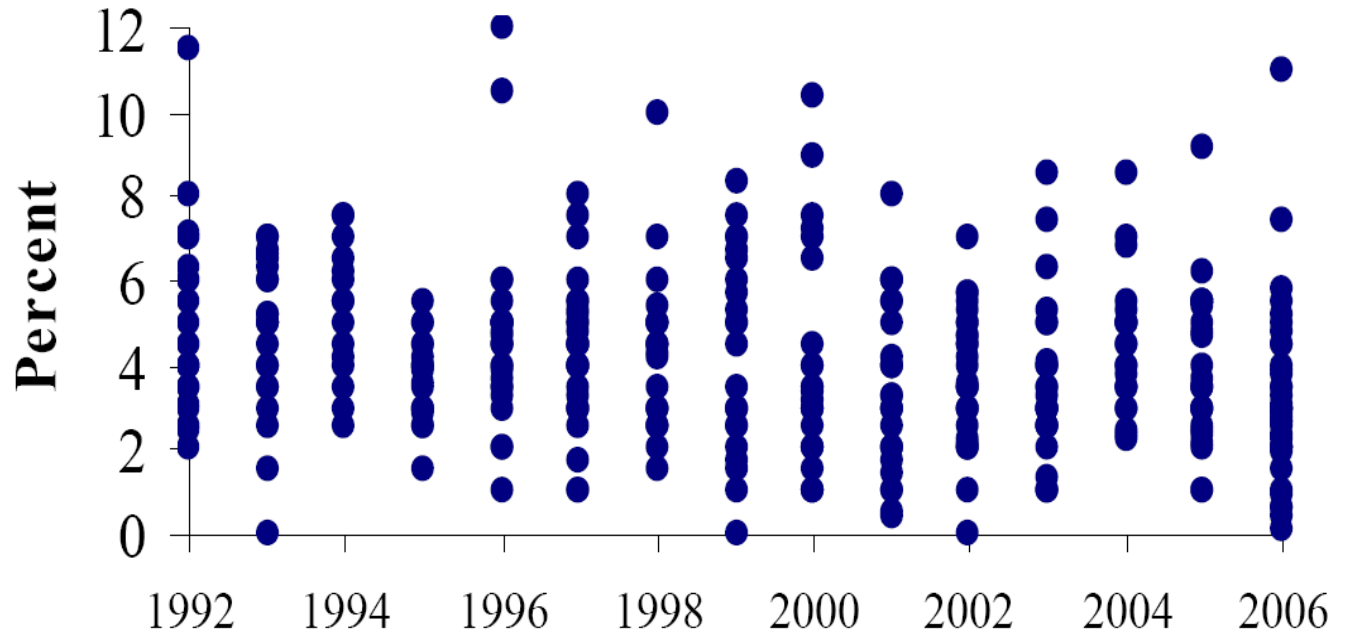


Figure 3

Term Structures of TIPS Yields

