

September 2016

You Get What You Pay For: Guaranteed Returns in Retirement Saving Accounts

William G. Gale, David C. John, and Bryan Kim

Gale and Kim: Brookings Institution and Retirement Security Project. John: Retirement Security Project

We gratefully acknowledge helpful comments from Jeffrey Brown, Hilary Gelfond, Ted Gayer, Joshua Gotbaum, Richard Kopcke, Aaron Krupkin, Olivia Mitchell, Kent Smetters, and Anthony Webb.

We thank the Laura and John Arnold Foundation for funding this work. The findings and conclusions contained within are solely the responsibility of the authors and do not necessarily reflect positions or policies of any organization with which they are associated or their funders

Introduction

The sharp downturn in the value of financial assets between 2007 and 2009 serves as a pointed example of how risky assets can quickly lose significant value. This experience, coupled with continuing concerns about retirement security, has generated new interest in the idea of having the government provide minimum rate-of-return guarantees for retirement savings accounts.

Guaranteed returns are not a new concept. Defined contributions plans in several countries provide minimum rate-of-return guarantees, as do some defined contributions plans in the United States. TIAA-CREF's "Traditional Annuity" provides a prominent example of an account with a guaranteed minimum return. Cash balance plans offer savers a fixed rate of return—thus, the guaranteed minimum return is equal to the ceiling on returns that the saver can receive. Many 401(k) or mutual funds offer "stable value" options that guarantee return of principal.

A variety of recent proposals would offer guarantees for new types of savings plans including some state-sponsored retirement savings plans for small businesses. State governments in California and Connecticut recently decided against requiring guarantees.

The key economic issues are the level of costs and benefits associated with a government-provided guarantee and who would bear the costs. Guarantees are a classic example of the economics dictum that it is impossible to get something for nothing. In principle, rate-of-return guarantees are simple: they could protect savers from losses and ensure that they receive at least a minimum return on their investments. In practice, they raise a variety of complex issues and are more costly than meets the eye. First, someone—the saver, the plan sponsor, or the taxpayers—has to pay for the guarantee. When the government pays the costs, budget documents tend to severely underreport the economic costs associated with the guarantees. Those costs are resources that have to be forgone in order to finance the promises. When private insurers offer guarantees, the costs, reflecting true economic costs more accurately, are often quite high. Second, the net benefits may not be as obvious as they seem, since markets often respond quickly and since for most people social security, Medicare, and housing are the source of the vast bulk of retirement resources.

Discussion of government provision of guarantees has often been confused by the comingling of different measures of costs. The expected, budgetary cost of a

guarantee is simply the discounted value of revenues and payments under the program, where discounting occurs at the government's risk-free rate. This measure describes the impact on the government's budget, but it is not sufficient to understand the economic costs and value associated with the program.

The economic costs are a measure of the value of the foregone resources used to implement the guarantee. This value is independent of whether the government or the private sector provides the guarantee. Ultimately, the level of economic costs and value associated with a guarantee depend on how high a rate-of-return is being guaranteed and what time period is covered. The allocation of those costs—to savers, plan sponsors, or taxpayers—depends on how the guarantee is financed. The financing could be done explicitly through general revenues or premiums paid by workers. It could be done implicitly via controls over the saver's portfolio choices, and or a cap on the maximum amount that the saver can earn, with profits above that level going to the government or insurer.

The rest of the paper is organized as follows. Section II discusses basic design elements. Section III discusses guarantees that exist in other countries and in certain United States plans, as well as recent policy proposals. Section IV discusses the determinants of the level and allocation of economic costs of the guarantees. Section V reviews recent estimates of these costs. Section VI provides concluding remarks.

Guarantee Design

A rate-of-return guarantee is essentially an insurance policy that ensures that a saver receives a certain return on his or her investments. When those investments earn less than the guarantee over a set time period, the saver receives the difference between the actual earnings and the promised amount from the guarantor. If the investments earn more than the guarantee, the investor receives the investment earnings; the insurer (government or private) does not make a payment.¹

While all minimum rate-of-return guarantees share these basic features, they can differ in a variety of ways. The guarantee can apply to investment returns in a particular year or to cumulative returns over a specified longer period. The promised rate of return can be constant over time or it can vary year-by-year in response to factors such as economic conditions. For example, an insurer might guarantee a minimum three percent return on contributions made in all years, or it might guarantee at least three percent on contributions made in the first year,

annual rate-of-return guarantee is 3 percent, the insurer commits to make up any shortfall between the actual account value at the end of the year and \$103.

¹ For those familiar with investing, the creation of a minimum rate-of-return guarantee is the equivalent of the saver/employee buying a put option from the insurer. A put option gives the saver the opportunity to sell an asset to the counterparty (the insurer), who is required to buy it if it is offered for sale by a given date for a given price. For example, if the saver contributes \$100 and the

but apply some other minimum, say 2.5 percent, on contributions made in following years. Also, the minimum guaranteed return might be enforced at the end of each specified time period or only when the employee changes jobs or retires.

Common rate-of-return guarantees include principal protection only (a guaranteed minimum nominal return of zero), a guarantee that the principal is returned with an adjustment for inflation (a guaranteed minimum real return of zero), or a guarantee based on the rate of return on a specific type of government bond or government bond portfolio. Other guarantees might be based on the rate of return on a specific market portfolio, sometimes expressed as a “reference portfolio” (Consiglio et. al 2015). Alternately, a guarantee might just promise a nominal return of a set level. Most nominal rate-of-return guarantees are in the 2-4 percent range. Some proposals also include protection against catastrophic market events by limiting losses to a set percentage of the initial investment.²

Guarantees are not free. They might be paid for explicitly, via insurance premiums that savers or plan sponsors pay. Alternatively, the costs may be implicit. For example, savers can pay for the guarantee by accepting restrictions on their investment portfolio or allowing the insurer to manage the fund and pay a minimum return plus any additional amount that trustees deem appropriate. In both of these cases with implicit payments, the costs take the form of the saver forgoing potentially higher returns on their investments.

Another way the saver could pay for a minimum guarantee is by selling some of the upside potential returns (Feldstein and Rangelova 2001a; Smetters 2002). In such a situation, the saver would be guaranteed a minimum rate of return, but there would be a ceiling on the maximum return he could keep from his investments, with any actual return above the ceiling going to the insurer. This combination is usually known as a “collar.”³ For example, the saver might be guaranteed that his investments would earn no less than three percent annually. In exchange, the saver would forfeit any upside beyond a specific ceiling (e.g., six percent annually) to the insurer. Hence, in this scenario, the saver’s portfolio is “collared” to generate only a 3-6 percent annual rate of return.

An appropriately designed collar allows the saver to receive a rate of return guarantee within a specified band and the insurer to be compensated at market rates for the

risk it is underwriting. Note that if the floor and the ceiling are the same rate of return, then the account simply has a guaranteed return, not just a guaranteed minimum return.

Figure 1 demonstrates the returns to buyers and sellers under minimum return guarantees and collars. With no guarantees or ceilings, the relationship between the rate of return earned on investments and the rate of return received by the saver/employee is given by the 45 degree line OBCD. The saver simply receives the actual return. If a guaranteed minimum return is put in place at level A, the returns to the saver as a function of actual returns are represented by ABCD (not including premium costs). This provides downside protection, as the saver is guaranteed a return of at least A. If there is a collar imposed at the actual rate of return given by point C, then the saver’s returns, as a function of actual returns, are shown by ABCE. (If the minimum guaranteed return and the maximum allowed return are the same – as, for example, in a cash balance plan – B and C would be at the same point and the line ABCE would be flat.)

Figure 1 also shows the net costs to the insurer as a function of actual returns received. In the graph, a positive number for the insurer reflects a positive cost, and a negative number represents a negative cost (i.e., a positive return). With a minimum guarantee, the returns to the insurer are provided by the line AFGH. When the asset generates actual rates of return below A, the insurer pays the difference between the actual return and the return given by A. When the asset generates actual returns at A or above, the insurer pays nothing. If a collar is in place at point C, the insurer’s costs are given by the line AFGI. For returns above C, the cost is negative; that is, the insurer receives the difference between the actual asset return and the ceiling on returns.

Existing and Proposed Guarantees

Minimum rate-of-return guarantees are offered in a number of existing and proposed plans in both the United States and a number of other countries around the world (Lachance et. al 2003; Turner and Rajnes 2003, 2009). As an example, several Latin American countries have instituted guarantees, often in conjunction with social security reforms. In Chile, the required return is defined in relation to returns in other plans. Uruguay offers a minimum guarantee that is the lesser of 2 percent real or average returns in the retirement system less 2 percent (Turner 2006).

² All of the options mentioned in the text guarantee rates of return on contributions. Rather than ensuring a minimum rate-of-return, guarantees could instead be provided for a minimum level of wealth at retirement, regardless of contribution levels. Neither rate-of-return guarantees nor minimum wealth guarantees, however, ensure a particular level of retirement income, since the rate that can be earned on an annuity varies over time. Thus, a third type of guarantee would ensure a minimum level of retirement income or minimum rate at which assets could be annuitized. In this paper, we focus on rate-of-return guarantees during the buildup stage.

³ The creation of a maximum rate-of-return allowed for the saver is the equivalent of the saver/employee selling a call option to the insurer. A call option gives the owner (the insurer, in this case) the opportunity to sell an asset to the counterparty (in this case, the saver/employee), who is required to buy it if it is offered for sale by a given date for a given price. For example, if the saver contributes \$100 and the annual rate-of-return maximum is 6 percent, the saver commits to give up any excess of the actual account value at the end of the year and \$106.

In Japan, defined contribution plans must offer at least one principal-guaranteed account. In Germany, Reister (DC) plans provide principal guarantees (Lachance and Mitchell 2003). Belgium provides a minimum return of 3.75 percent for employee contributions and 3.25 percent for employer contributions (Muir and Turner 2011). In Switzerland, pension funds must meet a minimum threshold return that the government sets and periodically adjusts (Muir and Turner 2011). Denmark has a nationwide, mandatory defined contribution plan that is required to provide a minimum return tied to current long-term interest rates (Muir and Turner 2011). Several other OECD countries have less extensive guarantee programs. The European Union's (EU) Third Directive on Life Assurances stipulates that a rate-of-return guarantee cannot exceed 60 percent of the rate of return on government bonds denominated in the relevant currency, gross of taxes (European Commission 2002). The National Provident Fund in New Zealand offers a guaranteed nominal return of 4 percent (National Provident 2014).

In the United States, defined contribution accounts with guaranteed minimum rates of return are rare. The federal Employee Retirement Income Security Act (ERISA) requires that all investment returns be used solely for the benefit of pension participants, with reasonable allowance to defray administrative costs. This makes it very difficult to develop reserve funds that could be used to smooth out actual returns and help meet a guaranteed return target. As a result, United States plans that offer guaranteed minimum returns typically exist outside the reach of ERISA. This includes plans for state government employees in Ohio and Indiana, as well as plans for public employees in the three Texas counties that seceded from Social Security in the early 1980s.

TIAA-CREF's "Traditional Annuity" offers a guaranteed minimum rate-of-return. The guarantee is set annually at the time of the contribution and is valid on contributions made in that year until retirement. The rate for new contributions is adjusted each year in conjunction with economic conditions and has recently varied between 1 and 3 percent. The TIAA Board of Trustees may also declare, on a year-to-year basis, additional rates of return for a specific year only, but they are not guaranteed for future years. TIAA has credited such additional amounts every year since 1948. The rate of return (the sum of the guaranteed minimum and the credited rate) averaged 8.16 percent per year between 1980 and 2007 (Biggs 2010).⁴ By way of comparison, the S&P 500 averaged a return of 12.86 percent, the Lehman Brothers United States Aggregate Bond Index returned 9.01 percent, and the 10-year Treasury Bond yielded an average return of 8.88 percent (Bloomberg 2015; Damodaran 2015).⁵

Cash balance plans are a hybrid form of pension. From the saver's perspective, they closely resemble retirement

savings plans, but in legal terms, they are defined benefit plans and are regulated as such. Cash balance plans provide notional accounts for their participants, and annually credit a return to each participant's notional account. The plans essentially have a guaranteed return, with both a minimum and maximum set at the same level. As defined benefit plans, cash balance plans are backed by pooled assets that are managed by trustees and can be allocated in part to a reserve fund in years with high returns to help cover the implicit guarantee in low-return years.

The important point is not just that TIAA's "traditional annuity" and cash balance plans provide guaranteed minimum rates of return, but that they finance this guarantee by imposing a fairly low ceiling on returns. This strategy compensates the plan sponsor for risk and controls costs. Savers in these plans receive guaranteed minimum returns, and thus avoid the downside possibilities, but pay for this guarantee by giving up the upside potential for higher returns.

There have been numerous proposals for minimum guaranteed rates of return in the United States. Feldstein and Samwick (2001) propose private accounts in Social Security with a real principal guarantee (an inflation-adjusted minimum return of zero). Feldstein and Rangelova (2001a) propose what they call "accumulated pension collars" on private retirement accounts as a way of ensuring that partial privatization of Social Security would not reduce benefits relative to current law.

Ghilarducci (2007) proposes a new system of retirement savings accounts managed by a government entity with a minimum guaranteed real return of 3 percent. Importantly, this proposal would set up a system like TIAA, described above, where trustees would build and manage a reserve fund and could, but would not have to, allocate additional rates of return to savers (see also Ghilarducci, Hiltonsmith, and Schmitz 2012).

Costs and Benefits of Guarantees

The benefits of guarantees depend on their effects on expected level and variability of savers' retirement wealth balances, savers' risk aversion, and the share of retirement wealth that is expected to come from the guaranteed account. The value of a guarantee will also depend on a host of psychological factors, including loss aversion on the downside and regret aversion on the upside. Moreover, guarantees may exploit money illusions on a real basis.

Expected Costs versus Economic Costs

Analysis of the costs of guarantees has often proven to be confusing because of a failure to distinguish the different methods through which costs are measured. In particular,

⁴ The Traditional Annuity is a "guaranteed benefit policy." The assets backing guaranteed benefit policies are not plan assets subject to ERISA and thus not subject to ERISA requirements regarding asset management.

⁵ Average returns/yields calculated as compound annual averages over 1986-2007.

summing up the budgetary costs and receipts that were recorded or would be recorded by a government entity that is running a guarantee program reflects the expected costs to the government. This is not equivalent to the economic cost of providing a guarantee. The economic cost is the value—to the saver and the insurer—of the resources devoted to meeting the guarantee. It includes both actual costs paid out and any gains that might have been lost if the saver did not have that guarantee or had a different type of guarantee. Insurers also face economic costs that represent the risk of having underfunded liabilities. The determinants of economic costs are further discussed in the next subsection.

Feldstein and Rangelova (2001b) provide an example that can be used to distinguish between expected cost and economic cost. In their example, a 45 year old worker contributes \$1,000 to an account and allocates 60 percent to the S&P 500 index and 40 percent to corporate bonds. The account makes a single payout at age 65. Assuming that the account earns the historical rates of returns for this portfolio, the expected value of the payout would be \$3,510.⁶ However, this payout amount is subject to uncertainty and the actual return could be higher or lower. Assuming perfect markets, perfect deployment of options pricing techniques, and certain market parameters, the authors show that the saver could buy, at no explicit cost, a collar that provided a minimum guaranteed payout of at least \$2,000 and a maximum of \$2,610.

The example implies there is an enormous difference between the expected and economic costs of a minimum guaranteed return for this saver's portfolio. The economic cost to the saver of having such a guarantee would be substantial. The saver has an expected payout of \$3,510 but can only insure (at no explicit cost) a payout in the range of \$2,000 to \$2,610. That is, the saver would have to give up about 26 percent of the expected final payout in order to be guaranteed that he will receive at least 57 percent of the expected final payment but no more than about 74 percent of that amount. To be clear, the saver might value the reduced range of uncertainty in returns more than the costs and so might choose to buy the collar in question. However, the saver is not getting something for nothing; he is forgoing considerable, yet uncertain upside benefits in exchange for a less substantial, but certain range of benefits.

This distinction between expected and economic costs explains much of the divergence in the literature, which seems to reach two broad, seemingly contradictory conclusions. One conclusion is that it would not have cost the government (or another insurer) much, if anything, to guarantee reasonably high minimum returns in the past (Munnell et. al 2009; Stubbs and Rhee 2012), based on ex post returns. The other conclusion is that providing guarantees can result in substantial costs for both an

insurer and a saver. These findings may seem even more contradictory when it is noted that the prospective studies base their analysis on asset returns patterns and economic conditions that are taken from historical data. But the discrepancy between these findings is, in fact, easy to explain. The retrospective analysis focuses on expected (average) costs to the insurers (which, in the cited papers, is reported as the budgetary costs to the government) of providing guarantees, whereas the prospective analysis examines the economic costs to generate cost estimates of guarantee provision.

Determinants of the Level of Economic Costs

The level of economic costs of providing a rate-of-return guarantee will depend on several factors and can vary enormously across different types of guarantees. The first factor is simply the level of the guarantee that is provided. Other things equal, the costs of providing principal guarantees (i.e., a zero nominal return) will be less than the cost of providing any level of positive nominal return. Likewise, as long as inflation is positive, ensuring a real return of “x” percent will cost more than ensuring a nominal return of “x” percent.

A second factor is the time horizon of the guarantee. This can work either way—a longer time horizon can increase or reduce the cost of guarantees depending on the interplay between the guarantee, the saver's portfolio, and the pattern of asset returns (Lachance and Mitchell 2002, 2003). Guarantees that are “tested” more often (e.g., a guarantee that is applied annually, as opposed to only at retirement or a job change) will be more expensive.

Lachance and Mitchell (2002) argue that the determination of the economic cost should not depend on whether the government or a private insurer provides the guarantee. As they write: “While alternative approaches might be valid depending on the use to which they are put, here we propose to define guarantee cost such that the values generated indicate true economic resource costs. If the guarantee commitment were made to the capital market investors, the value of the guarantee could be determined by using option pricing techniques and a market value approach for the guarantee could be derived. Since the nature of the counterparty should not influence the economic value of the liability, the approached uses with capital market investors should also be valid with...” government provision of insurance [emphasis added] (Lachance and Mitchell 2002, 7).

This suggests that, to a first-order approximation, the *economic* costs would be the same if the insurance were provided by the government or by the private sector. There would be various differences in actual pricing, of course. A guarantee set by the government might be priced with political economy factors in mind. On economic grounds, however, a guarantee set by the government should factor

⁶ In this case, the historical rates of return have a mean of 5.5 percent and standard deviation of 12.5 percent.

in the economy-wide marginal costs of funds. A guarantee set by the private sector would need to account for profits, administrative and regulatory costs, risk management, as well as any market imperfections.

In addition, the government may be able to handle certain long-lived risks better than the private sector, even abstracting from political economy considerations and the day-to-day costs of running a firm. As Smetters (2002) points out, the proposed guarantees that would cover very long time periods can be better handled by government, as the private market does not typically provide such lengthy guarantees. Savers need to know that the insurer will be able to stay in the market long enough to fulfill its contractual obligations. In the presence of non-diversifiable financial risk or intergenerational risk, government is probably better suited than the market to smooth the associated risks.

The private sector would likely either charge very high fees to compensate for taking on such risks, because otherwise it would be unable to provide insurance against massive investment risks, or charge a lower rate and create the potential risk of needing to be bailed out.

Informal evidence suggests that quotes offered from financial institutions for various guarantees are typically higher and often much higher than would be suggested by theoretical calculations using perfect markets. This discrepancy presumably reflects some imperfection in the private market. This issue of the potential existence of imperfections in the market for minimum rate-of-return guarantees is highlighted by the fact that financial markets routinely provide other types of guarantees—fixed annuities, life insurance, stable value funds, etc.

Determinants of the Allocation of Costs

As noted above, someone has to pay for the economic costs of the guarantee. Obviously, one option is for taxpayers to bear the burdens via general revenues. Another way to cover the costs would be for workers to pay premia. A third approach would impose the costs on savers by turning the unencumbered offer of a minimum guarantee into an offer that provides the minimum guarantee, but also gives the insurer a portion of the upside returns. The most obvious option in this regard is a collar. The cost to the insurer of providing a collar is lower than the cost of providing the same minimum guarantee without a ceiling. The saver would pay for this feature by forgoing returns above a certain level.

Likewise, allowing the insurer to use some of the actual returns from the saver's portfolio in excess of the guaranteed rate to create a reserve fund that can be used to supplement actual returns in years when returns are lower than the guaranteed rates would shift costs to the saver. Both the TIAA traditional annuity and cash balance

plans are examples of this mechanism, and Ghilarducci (2007) includes this feature as a central part of her proposal for guaranteed returns. To be clear, this does not reduce the overall economic costs of the guarantee, it just provides a way for the saver to compensate the guarantor for taking on risk.

More general, restrictions on the savers' portfolio composition impose costs on savers. Actual portfolio composition is determined by the restrictions placed on portfolio contents by the insurer and the saver's subsequent portfolio construction given those restrictions. It is well understood that, with a minimum rate-of-return guarantee, a saver has an incentive to pursue more risky returns, since the guarantee protects the saver from downside risk. What appears to attract less attention, however, is the notion that portfolio restrictions can materially impact the risk associated with minimum guarantees. As an extreme example, a guarantee of principal repayment can be honored at zero risk to the insurer by requiring that the saver invest his entire portfolio in FDIC-insured bank accounts. As long as each bank account holds less than the maximum FDIC guarantee, there is no risk of loss. Likewise, a minimum guarantee of the return on Treasury bonds or a broad stock index can be provided by an insurer at no cost provided that the saver is required to invest his entire portfolio in Treasury bonds or the broad stock index in question. These portfolio restrictions act by exactly matching the risks associated with the guarantee and the risks associated with the assets backing the guarantee. By doing so, they entirely eliminate the risk of insuring the restricted portfolio, and of course, at the same time, they eliminate any benefit of insuring the restricted portfolio. As with collars, tight portfolio restrictions do not eliminate or even affect the total economic costs. They just provide a way for savers to bear the costs.

The minimum guarantee imposes economic costs on the insurer only when there are different risks embodied in the minimum guarantee and in the saver's portfolio. For example, a minimum guaranteed that is tied to inflation or interest rates creates risk for the insurer to the extent that saver's portfolio contains investments like stocks.

These conclusions imply that any guaranteed return below the risk-free rate of return can be provided at zero risk to the insurer *if suitable portfolio restrictions are placed on the saver*.⁷ Likewise, even returns that are *expected* to be above the risk-free rate – such as that on a broad portfolio of stocks – can be insured at no or little cost if the saver's portfolio is required to match the nature of the guarantee. But “bond style” guarantees (that insure something akin to an interest rate) are risky to the insurer and are of benefit to the saver if they underwrite portfolios that contain riskier assets like stocks; as expected, the risk to insurers and benefits to the saver grow with the equity share of the portfolio.

⁷ The saver may still be required to pay fees that cover administrative or regulatory costs.

Previous Estimates

In this section, we review several recent papers that estimate the costs of rate-of-return guarantees under different economic conditions and assumptions. Cost estimates differ widely depending on the assumptions used. Earlier work on the cost of pension guarantees includes Pesando (1982), Marcus (1985, 1987), Bodie and Merton (1993), Bodie (2001b), Smetters (2001, 2002) and Feldstein and Rangelova (2001).

Munnell, et. al (2009) provides both retrospective and prospective estimates of the costs of guarantees. On a retrospective basis, they find that, under certain circumstances and assumptions, including the use of ex post returns, the *budgetary (expected)* cost to the government of having provided even a fairly high minimum rate-of-return guarantee would have been quite small. On a prospective basis, they find that the *economic* costs of providing guarantees going forward could be expensive under a number of conditions, even under assumptions that would generate the same low budgetary costs that they obtained using historical data. But they also show that under a variety of conditions, the economic costs of providing certain guarantees would be small or vanishing.

Their retrospective calculations are based on a model where workers enter the workforce at age 22, work for 43 years with real annual wage growth of 2 percent, and then retire at age 65. Workers contribute 4 percent of their wages to a retirement account and invest the funds entirely in equities. The calculations cover 84 cohorts, beginning with the cohort of 24-year olds in 1883, who reached age 65 in 1925, and ending with the cohort that was 65 years old in 2008. During that period, real stock market returns averaged 7.6 percent per year, with a standard deviation of 19.5 percent.

The authors examine the frequency and size of payments that a guarantor would have had to make in order to provide workers when they reached age 65 with real returns ranging from 2 percent to 6 percent on their lifetime contribution. Remarkably, they show that guaranteeing a real return of 3 percent on lifetime contributions would never have induced the insurer to have to make a payment. Moving to a real return of 4 percent would have triggered payments only three times, twice in the onset of the Depression—1931 and 1932—at 0.06 percent and 0.17 percent of GDP, and once in 1941, at just 0.01 percent of GDP.⁸ Even a real required 5 percent rate of return would have required payment in only eight years, all between 1931 and 1948, and averaging 0.35 percent of GDP in those years (and averaging just 0.13 percent other than 1931- 1932). However, guaranteeing a real return of six percent would have cost substantially more – payments would have been triggered in 27 of 84 years and would have averaged 0.57 percent of GDP in those years.

These calculations assume that the insurer only provides a floor on returns. The net cost of guarantees to the insurer, however, depends on whether there is a ceiling as well. If there were a guaranteed real annual return (i.e., both a floor and a ceiling) of 6 percent on lifetime contributions, with the insurer being able to keep the upside returns as reserves to be used when needed – the collar would have had negative net costs of 0.8 percent of GDP over the entire time period. It is crucial to note this retrospective analysis looks at the insurer's (presumably the government's) budget costs – the expected values of payments and receipts – it does not calculate the economic cost of providing the insurance. Nevertheless, the analysis shows that the expected budgetary costs could have been negative over the 1925-2008 period even under a very generous guarantee program.

Moving to analysis of the prospective economic costs to the insurer, the authors use finance theory to calculate the market price of future guarantees. Their estimates of the prospective economic costs are based on several key assumptions: the insurers' aversion to risk matches that of the market; the saver is invested in an all-equity portfolio; future equity returns follow a random walk with the same mean and standard deviation as the historical data; and there is a 2 percent real risk-free rate of return.

Table 1 shows their estimates. The price of a floor at 2 percent real (the risk-free rate) is 29 percent of contributions. That is, for every dollar contributed to the plan, an additional 29 cents would have to be paid just for the guarantee. Thus, a saver would have to contribute \$1.29 to make a \$1 contribution with a guaranteed return. Another way to think of this is that for every dollar the saver contributes, 22.5 cents goes to a guarantee, and only the remaining 77.5 percent is invested. This estimate depends critically on the saver's portfolio. If the saver were required to invest in the risk-free asset, the costs would be much lower. Any other administrative costs for the plan would be on top of the cost for the guarantee.

As the floor rises, the required fee rises more than proportionately. A 3 percent real floor would cost 46 percent of contributions, while a 4 percent real floor would cost 71 percent of contributions, and a 6 percent real floor would cost 157 percent of contributions. Obviously, such premiums would make it virtually impossible for any saver to achieve any rate of return in excess of the guarantee rate even with a portfolio of 100 percent equities.

As before, if there is a ceiling as well as a floor, the economic costs to the insurer and the saver fall. Guaranteeing a fixed—minimum and maximum—return of 2 percent would be costless to the insurer under the stated assumptions (since it is the risk-free rate). Guaranteeing 4 percent would require a payment of 55 percent of contributions. More interesting results appear if one allows the floor and the ceiling to diverge. For example, a 2

⁸ While a cost in terms of percent of GDP might be sustainable for a government, it would be far higher than any private insurer could bear.

percent floor with a 6 percent ceiling would cost 22 percent of contributions.

If the insurer is only half as risk-averse as the market, these figures change substantially and the economic costs fall dramatically. Under these assumptions, the price of a 2 percent floor is only 13 percent. It can even be coupled with a 7 percent ceiling and still have a negative economic cost to the insurer. Indeed, a floor/ceiling combination of 3 percent and 7 percent or 4 percent and 6 percent could be handled at no economic cost to the insurer. Interpolating from the table results, a guaranteed real return of about 4.6 percent should be possible to achieve at no net economic cost to the insurer under these assumptions. Companies offering such a guarantee would still need to charge for administrative costs associated with the guarantee and would include an amount for profit.

Lachance and Mitchell (2002, 2003) estimate the costs of minimum guarantees for principal repayment (that is, a guaranteed minimum nominal return of zero) and one equal to the Treasury 10-year bond rates. They define the cost as the price that would be charged by market participants to provide the guarantee, using option pricing techniques, and obtain several major results.

First, as expected, the more stringent the minimum guarantee, the more expensive it is to finance. If backed by a saver portfolio that is split evenly between stocks and bonds, nominal and principal guarantees have almost no cost – less than 0.5 percent of contributions over 40 years. With the same portfolio, bond return guarantees, which are significantly more stringent than the other principal guarantees, cost more, 16.1 percent of contributions.

Second, the costs of insurance depend dramatically on the saver's portfolio. A higher equity share in the saver's portfolio raises the cost. This occurs because the higher expected return of equities becomes more volatile over the long-term as the equity share rises (Bodie 2001a). In contrast, if the saver can be restricted to holding an all-bond portfolio, or chooses to do so voluntarily, there is no cost to any of the guarantees examined.

Third, the results vary in a non-monotonic way with respect to the duration of the guarantee. The longer time duration raises the costs of the bond return guarantee but reduces the cost of the nominal principal guarantee. For example, for an all-stock portfolio, the cost of a nominal principal guarantee falls from 3.6 percent of lifetime contribution over 10 years, to just 0.8 over 40 years. In contrast, for the same all-stock portfolio, the cost of a bond return rises from 16.1 percent of contributions over ten years to 31.3 percent over 40 years. In both cases, the value of the underlying stock portfolio is becoming more volatile over time, which should raise guarantee costs (Bodie 2001a). But in the case of the principal guarantee, the real value of

the guarantee decays over time, assuming inflation is positive, which offsets the volatility effect.

Lachance and Mitchell (2002) also report data on required fees as a share of net asset value, shown in the right half of Table 2. Two of the three main results above continue to hold. More stringent guarantees generate higher costs, controlling for the saver's portfolio. Higher equity shares raise costs as well. What is different is that longer duration uniformly reduces costs as a share of net asset value, including for the all-equity proposal, whereas costs rose with the duration for the all-equity portfolio as a share of lifetime contributions. This difference occurs because expected net asset value is higher when savers have an all-equity portfolio than when they have a portfolio evenly split between stocks and bonds, but lifetime contributions are the same in the two cases.⁹

Grande and Visco (2010) examine the potential economic costs of insuring minimum rate-of-return guarantees in the European market. They examine 3 minimum guarantees: a nominal principal guarantee (a zero nominal rate of return), a 2.5 percent nominal (which they take as approximately a real principal guarantee -that is a minimum real return of zero) and a return equal to the annual nominal GDP growth rate. Table 3 shows that their estimates of the economic costs of guarantees share several features with Lachance and Mitchell (2002, 2003), but differ in an important way as well. Grande and Visco (2010) obtain three results very similar to Lachance and Mitchell: 1) when measuring costs relative to lifetime contributions, the costs are higher for more stringent guarantees; 2) the costs rise with the equity share of the saver's portfolio and are zero for the nominal and real principal guarantee if the saver invests the portfolio entirely in bonds; and 3) the costs of the most stringent portfolio option rise over time with the all-equity proposal. As with Lachance and Mitchell (2002), the costs of the most stringent guarantee fall over time as a share of net asset value, even if the saver holds all equities in his portfolio. Again, this is due to the expected growth over time of the assets in the portfolio.

Grande and Visco (2010) have significantly higher cost estimates for rate of return guarantees than those in Lachance and Mitchell (2002). Grande and Visco (2010) estimate the cost of a nominal principal guarantee as 11.4 percent of contributions over 10 years and 6.39 percent of contributions over 40 years when the saver has an all-equity portfolio. Lachance and Mitchell (2002, 2003) estimate these costs at 3.6 percent and 0.8 percent, respectively. Similar differences persist across different saver's portfolios and different types of guarantees. These differences stem from the fact that the two sets of authors use different assumptions about risk-free interest rates, stock volatility, and the source of the error terms used to simulate randomized stock market time series patterns.¹⁰ A key factor in the increased cost estimates by Grande and Visco

⁹ Because even though equities are assumed to be more volatile, they are assumed to have a greater average return than bonds.

¹⁰ First, Grande and Visco use a higher risk-free interest rate. Second, they assume that stock market prices follow a stochastic process with jumps. With

no jumps, the underlying process has annual volatility of 16 percent. The process has an expected 1.8 jumps per year, with an average jump size of -12.8 percent (i.e., the stock market declines). In contrast, Lachance and Mitchell assume volatility of 20 percent with no jumps. Third, Grande and Visco generate the

(2010) is their assumption that stocks will be much more volatile, which increases the possibility that the guarantor will suffer losses.

Interestingly, Grande and Visco (2010) also provide estimates of the costs of guarantees assuming that savers use a life-cycle investment strategy in which they hold all of their portfolio in equities until age 55 and then shift the percentage linearly to zero by age 65. The benefits of the life-cycle portfolio strategy are clear. The 40-year costs of guaranteeing a fund are lower for all three guarantee options under the life-cycle investment strategy than they are under either the all-equity portfolio or the 50/50 split portfolio.

Scheuenstuhl et. al (2011) examine several different guarantees, including a nominal return of zero, a real return of zero, nominal returns of 2 and 4 percent, and a floating guarantee of a return that equals current government bond interest rates plus one percentage point. For their calculations, the authors assume investors use a life-cycle investment strategy over 40 years using a strategy similar to Grande and Visco (2010). In this paper, the saver's portfolio is 80 percent invested in equities and 20 percent in government bonds until age 55. At that point, the equity allocation is linearly reduced to 20 percent by age 65. The authors then specify how much money would be needed to finance a given guarantee in the private market; that is they calculate the fair value of a guarantee based on a risk-neutral pricing framework. However, they also note that guarantee prices in the real world might deviate from their estimates because of transaction costs, liquidity premia, and solvency and capital requirements, and thus could be higher than those in the paper.

The general results of Scheuenstuhl et. al (2011) follow those of the other papers. As in the studies described above, the authors find that costs rise with the stringency of the guarantee. They also find that costs of providing a guarantee fall as the duration of the guarantee period increases. Given that the paper assumes that savers follow a life-cycle investment strategy, its results are most comparable to, and are consistent with, those in Grande and Visco (2010), which also assumes that savers are life-cycle investors. The paper's cost estimates differ from those of other papers that use different portfolio assumptions. This highlights the importance of the saver's portfolio (either through the saver's choice or through restrictions placed by the insurer) in determining the costs of a guarantee.

The importance of how the portfolio is structured is emphasized by additional life-cycle investment strategies that start with lower initial allocations to equity (50 percent and 20 percent, rather than the 80 percent in the base

case), and have lower costs. This is consistent with the general point that other things equal (where other things might include the final allocation to equity and the date at which a saver starts shifting to bonds), a higher equity share in saver's portfolio raises the costs of providing a minimum guarantee.

Finally, although the cost estimates are not quite comparable to other research because they use different saver portfolio assumptions, the estimates in Scheuenstuhl et. al (2011) appear to be between those of Lachance and Mitchell (2002, 2003) and Grande and Visco (2010). For example, the paper estimates that guaranteeing a nominal return of zero over 40 years in their life-style portfolio would require premiums of 1.244 percent of contributions. Grande and Visco (2010) find that the same guarantee would vary in cost depending on whether savers' portfolios were all-equity (6.39 percent), all bonds (0.48 percent), or invested in a life-cycle portfolio (0.22 percent). Lachance and Mitchell (2002, 2003) find that the same guarantee would cost 0.8 percent of contributions if savers held an all-equity portfolio and nothing if savers held evenly split portfolios.¹¹


Conclusion

The steep losses suffered by savers close to retirement during the recent financial crisis have motivated an increase in attention to rate-of-return guarantees for retirement saving plans. While guarantees in various forms clearly offer some benefits to savers, the benefits come at a cost. The costs can be paid in many different ways, including insurance premiums, caps on the maximum returns that savers can receive on their investments, or portfolio restrictions. The last option may also serve to cap returns and limit the risks that savers can take. In any of those cases, the true economic costs of providing the guarantee will substantially exceed the expected budgetary costs to the government of offering the guarantee. A private insurer would likely charge the economic cost to offer a guarantee. The government may not, for political reasons, but that does not make the economic costs disappear.

distribution of shocks to stock returns using data from 2000-2002. Lachance and Mitchell use stock data from 1926-2000. The first difference should reduce Grande and Visco's estimates relative to Lachance and Mitchell's, while the second should have the opposite effect. The impact of the third difference is unclear a priori.

¹¹ Scheuenstuhl et. al (2011) also utilize risk-neutral valuation techniques to derive their cost estimates. However, Scheuenstuhl et al. conceptualize bond returns using a Hull-White process rather than the Vasicek process that Lachance and Mitchell (2002, 2003) and Grande and Visco (2010) employ. The primary difference between the Vasicek and Hull-White processes is that the Hull-White model allows for time dependency in some of the coefficients.

Figure 1
Returns to Savers and Insurers under Various Guarantees



Source: Author's Calculations

Table 1

Cost Estimates of Alternative Minimum Guarantees from Munnell et al. 2009¹

(As a Percentage of Contributions)

Rate of Return	Insurers are as Risk Averse as the Market			Insurers Half as Risk Averse as the Market		
	Price of Floor	Price of Ceiling	Price of Collar ²	Price of Floor	Price of Ceiling	Price of Collar ²
2%	29	29	0	13	97	-84
3%	46	22	24	23	83	-60
4%	71	16	55	40	68	-28
5%	107	11	97	66	53	13
6%	157	7	150	106	40	66
7%	224	4	220	163	28	135

¹ Saver's portfolio assumed to be all equity.

² Price of Collar = Price of floor less Price of ceiling

Source: Munnell et al. (2009)

Table 2

Cost Estimates of Alternative Minimum Guarantees from Lachance and Mitchell (2002, 2003)

Saver's Portfolio	Investment Horizon	Guarantee Cost as a Percentage of Lifetime Contributions			Guarantee Cost as a Percentage of Net Asset Value		
		Principal (Zero Nominal Return)	Real Principal (Zero Real Return)	Treasury 10-year Bond Rate	Principal (Zero Nominal Return)	Real Principal (Zero Real Return)	Treasury 10-year Bond Rate
100% Equity	10 years	3.6	8.2	16.1	0.60	1.36	2.67
	40 years	0.8	5.9	31.3	0.03	0.24	1.27
50% Equity/ 50% Bonds	10 years	0.2	2.0	8.1	0.04	0.33	1.35
	40 years	0.0	0.5	16.1	0.00	0.02	0.65
100% Bonds	10 years	0.0	0.0	0.0	0.00	0.00	0.00
	40 years	0.0	0.0	0.0	0.00	0.00	0.00

Source: Lachance and Mitchell (2002, 2003)

References

- Aresimowicz, Joe and Martin Looney. (2014). S.B. No. 249. *An Act Promoting Retirement Savings*.
- Biggs, Andrew G. (2010). "Improving the Social Security Statement." *SSRN Electronic Journal*. doi:10.2139/ssrn.1991397.
- Bloomberg LP. (2015). Total Return Analysis graph for Barclays Aggregate US Bond Index. via *Bloomberg LP*.
- Bodie, Zvi. (2001a). "On The Risk of Stocks in the Long Run." *SSRN Electronic Journal*. doi:10.2139/ssrn.5771.
- Bodie, Zvi. (2001b). "Retirement Investing: A New Approach." *SSRN Electronic Journal*. doi:10.2139/ssrn.260628.
- Bodie, Zvi and Robert C. Merton. (1993). "Pension Benefit Guarantees in the United States: A Functional Analysis." In *The Future of Pensions in the United States*. University of Pennsylvania: University of Pennsylvania Press.
- Consiglio, Andrea, Michele Tumminello, and Stavros A. Zenios. (2015). "Designing Guarantee Options in Defined Contribution Pension Plans." *SSRN Electronic Journal*. doi:10.2139/ssrn.2563087.
- Damodaran, Aswath. (2015). *Annual Returns on Stock, T.Bonds and T.Bills: 1928 – Current*. NYU Stern Business School.
- De León, Kevin. (2012). S.B. No. 1234. *California Secure Choice Retirement Savings Trust Act*.
- European Commission. (2002). *Directive 2002/83/EC concerning life assurance*.
- Feldstein, Martin and Elena Ranguelova. (2001a). *Accumulated Pension Collars: A Market Approach to Reducing the Risk of Investment-Based Social Security Reform*. In *Tax Policy and the Economy, Volume 15*, edited by James M. Poterba. MIT Press.
- Feldstein, Martin and Elena Ranguelova. (2001b). *Individual Risk in an Investment-Based Social Security System*. w8074. Cambridge, MA: National Bureau of Economic Research. <http://www.nber.org/papers/w8074.pdf>.
- Feldstein, Martin and Andrew Samwick. (2001). *Potential Paths of Social Security Reform*. w8592. Cambridge, MA: National Bureau of Economic Research. <http://www.nber.org/papers/w8592.pdf>.
- Ghilarducci, Teresa. (2007). *Guaranteed Retirement Accounts: Toward Retirement Income Security*. EPI Briefing Paper 204. Economic Policy Institute.
- Ghilarducci, Teresa, Robert Hiltonsmith, and Lauren Schmitz. (2012). *State Guaranteed Retirement Accounts: A Low-Cost, Secure Solution to America's Retirement Crisis*. The New School: Schwartz Center for Economic Policy Analysis.
- Grande, Giuseppe and Ignazio Visco. (2010). *A Public Guarantee of a Minimum Return to Defined Contribution Pension Scheme Members*. Working Paper 742. Banca D'Italia.
- Lachance, Marie-Eve and Olivia Mitchell. (2002). *Understanding Individual Account Guarantees*. w9195. Cambridge, MA: National Bureau of Economic Research. <http://www.nber.org/papers/w9195.pdf>.
- Lachance, Marie-Eve and Olivia S Mitchell. (2003). "Guaranteeing Individual Accounts." *American Economic Review* 93 (2): 257–60.
- Lachance, Marie-Eve, Olivia S. Mitchell, and Kent Smetters. (2003). "Guaranteeing Defined Contribution Pensions: The Option to Buy Back a Defined Benefit Promise." *Journal of Risk & Insurance* 70 (1): 1–16. doi:10.1111/1539-6975.00044.
- Marcus, Alan. (1987). *Corporate Pension Policy and the Value of PBGC Insurance*. In *Issues in Pension Economics*, edited by Zvi Bodie, John B. Shoven, and David A. Wise. University of Chicago Press.
- Marcus, Alan J. (1985). "Spinoff/Terminations and the Value of Pension Insurance." *The Journal of Finance* 40 (3): 911–24. doi:10.1111/j.1540-6261.1985.tb05018.x.
- Muir, Dana M. and John A. Turner. 2011. *Imagining the Ideal Pension System: International Perspectives*. W.E. Upjohn Institute for Employment Research: Upjohn Institute Press
- Munnell, Alicia H. Alex Golub-Sass, Richard W. Kopcke, and Anthony Webb. (2009). *What Does It Cost to Guarantee Returns?* 9-4. Boston College: Center for Retirement Research.
- National Provident. (2014). *National Provident Pension Scheme: Summarized Annual Report*.
- Pesando, James. (1982). *Valuing Pensions (Annuities) with Different Types of Inflation Protection in Total Compensation Comparisons*. NBER Working Paper

956. National Bureau of Economic Research.

Scheuenstuhl, Gerhard, Sandra Blome, David Karim, Matthias Moch, and Stefanie Brandt. (2011). *Assessing the Nature of Investment Guarantees in Defined Contribution Pension Plans*. Financial Study. Risklab.

Smetters, Kent. (2001). *The Equivalence of the Social Security's Trust Fund Portfolio Allocation and Capital Income Tax Policy*. w8259. Cambridge, MA: National Bureau of Economic Research. <http://www.nber.org/papers/w8259.pdf>.

Smetters, Kent. (2002). "Controlling the Cost of Minimum Benefit Guarantees in Public Pension Conversions." *Journal of Pension Economics and Finance* 1 (01). doi:10.1017/S1474747202001014.

Stubbs, David M. and Nari Rhee. (2012). "Can a Publicly Sponsored Retirement Plan for Private Sector Workers Guarantee Benefits at No Risk to the State?" Policy Brief. UC Berkeley: Center for Labor Research and Education.

Turner, John A. (2006). *Individual Accounts for Social Security Reform: International Perspectives on the U.S. Debate*. W.E. Upjohn Institute for Employment Research: Upjohn Institute Press.

Turner, John A. and David Rajnes. (2003). "Retirement Guarantees in Voluntary Defined Contribution Systems." In *The Pension Challenge*, edited by Olivia S. Mitchell and Kent Smetters. Oxford: Oxford University Press.

Turner, John A. and David Rajnes. (2009). "Guarantee Durability: Pension Rate of Return Guarantees in a Market Meltdown." Collegio Carlo Alberto: Center for Research on Pensions and Welfare Policy.