Economic Studies at BROOKINGS

September 2016

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

Policy Brief[1]

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

1: This policy brief is based on a Retirement Security Project working paper of the same title, available at www.brookings.edu/project/retirement-security-project.

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.

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.

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 that the guarantee, the investor receives the investment earnings; the insurer (government or private) does not make a payment.

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, 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.

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 Ranguelova 2001, 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." 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.

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). 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, US 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). By way of comparison, the S&P 500 averaged a return of 12.86 percent, the Lehman Brothers U.S. Aggregate Bond Index returned 9.01 percent, and the 10-year Treasury Bond yielded an average return of 8.88 percent.

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 Ranguelova (2001) 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 three 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, 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.

Determinants of the 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. 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 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.

More general, restrictions on the savers' portfolio composition impose costs on savers. 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.

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 or private insurer 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.

References

Biggs, Andrew G. 2010. “Improving the Social Security Statement.” SSRN Electronic Journal. doi:10.2139/ssrn.1991397.

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.

European Commission. 2002. Directive 2002/83/EC concerning life assurance.

Feldstein, Martin, and Elena Ranguelova. 2001. 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 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.

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.

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

National Provident. 2014. National Provident Pension Scheme: Summarized Annual Report.

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.

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 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.