Pension guarantees in developing countries

Ajay Shah*
ajayshah@mayin.org
http://www.mayin.org/ajayshah

July 7, 2006

Abstract

*The views in this paper are those of the author and not his employers. I am grateful to Mike Orszag, Robert Palacios and Renuka Sane for useful discussions.

Contents

1 Motivation **3**

2 Understanding guarantees **4**
2.1 Alternative pension guarantee structures 4
2.2 Summary statistics about the cost and benefit of a guarantee 5

3 "Cassandra" - a software system for analysing guarantees **6**
3.1 The model 6
3.2 Parameters about a country 8
3.3 Generic parameters 10
3.4 Software implementation issues 10
3.5 Example: India 11
3.6 Limitations 17

4 Policy issues of operationalising guarantees **19**
4.1 The fiscal dimension 19
4.2 Implementation difficulties in the real world 20
4.3 Feasible policy alternatives 20

5 Conclusion **21**

1 Motivation

In recent decades, pension systems have increasingly moved towards individual-account, defined-contribution programs. The equity premium (Siegel, 2002) has come to prominence in thinking about investment policies followed by pension funds. Investment policies followed by pension funds have come to increasingly harness risk premia by exposing portfolios to systematic risk factors.

A major difficulty with DC pension systems is the problem of investment risk which has to be borne by the participant. The wealth at retirement date becomes uncertain in a DC system, as does the stream of annuities, faced upon retirement. This exposure to uncertainty reduces the welfare of risk-averse workers.

This has generated heightened interest in the questions of risk management for pension investments. There are many mechanisms of risk management which can be adopted. At the simplest, participants can voluntarily modify their exposure to high-volatility asset factors depending upon their age profile and their risk preferences. Such individual risk management efforts could be supplemented by *guarantees*, where there may be a role for public policy.

Sometimes, pension guarantees constitute contingent liabilities for taxpayers, and are a transfer from taxpayers to pension system participants. Alternatively, pension guarantees can be created without requiring support from the exchequer, by offering some reduction in the probability of low welfare outcomes, in return for lower expected returns. In formal terms, the stochastic environment induces a density function of pension wealth $g(v_T)$ at the terminal date T. Financial engineering defines a large range of alternative densities $h_1(v_T), \ldots h_N(v_T)$ which are attainable in the sense that a self-financing trading strategy can be designed which induces the density $h_i(v_T)$. The quest for pension guarantees enhances welfare when there is a self-financing but modified density $h_i(v_T)$ which yields a higher expected utility when compared with $g(v_T)$.

In recent years, there has been a heightened awareness about the "mistakes" made by participants faced with choice in a DC pensions ystem. One element of this is the unwillingness of unsophisticated workers to expose themselves to risk factors when faced with investment choices in a DC pension system. Such poor decision making hinders the accumulation of pension wealth. The introduction of certain kinds of guarantees might help to produce better decisions on the part of households about exposing the pension portfolio to systematic risk.

This paper seeks to offer new ideas on the valuation and consequences of pension guarantees, particularly in developing countries. In order to do this, we will explicitly articulate certain specific guarantee structures, and seek to obtain numerical values for the price profile of these guarantees in an idealised complete-markets setting, under certain assumptions about the parameters which characterise price processes.

This paper is not about the question of *whether* there is a role for equity investment in pension fund management. The goal of the paper is to obtain a better understanding the *consequences* of envisioning guarantee structures in a world where pension funds invest in the equity index.

2 Understanding guarantees

2.1 Alternative pension guarantee structures

There are dozens of different guarantee structures found in the real world. However, five kinds of alternative designs appear to be unusually important (Sin, 2002; Walliser, 2002).

0 *No guarantee.*

1 *Guarantee that avoids poverty.* In this structure, if the worker finds himself in poverty at the date of retirement, the pension system tops off the terminal pension assets in order to ensure that the annuity which is purchased is above the poverty line.

2 *Minimum 50% replacement rate.* In this structure, if the replacement rate drops below 50%, the pension system will top off terminal pension assets so as to reach a replacement rate of 50%.

3 *Guarantee that returns will beat inflation.* In this structure, if the real rate of return obtained by the participant on his entire lifetime of pension contributions drops below 0%, the pension system will top off terminal pension assets so as to atleast reach this 0% level. While this may also be called a "preservation of capital guarantee", it is stronger than the conventional interpretation of capital–preservation guarantees in that the guarantee is invoked when returns drop below 0% in *real* terms.

4 *Relative returns guarantee* If the participant in a DC pension system chose a fund manager who underperformed the average returns of the system, then the pension account is topped off to reach the average returns of the system.

In this paper, we analyse structures 0, 1, 2 and 3. The fourth structure - the relative returns guarantee - is much harder to analyse, since it is primarily about the difficulties of individual fund managers and not about investment risk associated with alternative asset classes.

Going beyond these relatively well-known guarantee structures, many creative modifications come about in real world pension systems, particularly given the complex processes of negotiation and incremental changes which take place in the public policy debates about pension reform. However, a thorough understanding of these four alternatives would constitute an important contribution into the public discourse on pension guarantees.

2.2 Summary statistics about the cost and benefit of a guarantee

When DC pensions are part of an existing or proposed pension system, the policy discourse is likely to touch upon the question of guarantees. At the simplest, politicians and participants who are uncomfortable with the volatility of financial markets might propose State-supplied guarantees, purely as a way of overcoming their mistrust of markets. At the same time, sophisticated proposals involving guarantees are also likely to be in the marketplace of ideas.

Across a range of possible guarantee structures, the discussion about guarantees would be illuminated if there was a better understanding about the cost and benefits of alternative guarantee structures. At an abstract level, if the density of pension wealth at date T under a given investment strategy is $g(v_T)$, a guarantee structure i induces a modified density $h_i(v_T)$. Policy analysis could ask which of the two distributions is associated with higher utility. However, such a formal analysis is not tractable in the real world.

At a practical level, it appears useful to summarise the costs and benefits of guarantees, in the context of an investment strategy, through four summary statistics:

- *The cost of the guarantee.* Guarantees are a complex contingent liability for public finance, with potentially enormous payments that could have to be made many decades in the future. All too often, policy discussions focus on the terms and conditions of the guarantee. In addition to this, the policy discourse needs to be informed by estimates of the price of alternative guarantee structures.
- *Probability of being in poverty.* The political rationale for guarantees is often expressed as an urge to eliminate the possibility that a pension system participant finds himself in old age in a state of poverty. This motivates measurement of the probability of arriving in old age and ending up with an annuity flow which is below poverty.
- *The median replacement rate.* Random outcomes associated with risky

asset price outcomes induce ex-ante uncertainty about the replacement rate. The median replacement rate serves as a useful summary statistic about the distribution of the replacement rate.

- *Investment risk as measured by the interquartile range of the replacement rate.* One useful measure of investment risk is the inter-quartile range of the replacement rate. This can be computed for various kinds of investment strategies involving no guarantees, and the impact of a guarantee on this measure of risk can then be measured.

The comprehension and thus the utilisation of these four numbers in realworld policy discourse involving nonspecialist participants appears to be feasible and useful.

Each of these four numbers is sensitive to two parameters: the equity exposure, and the initial income of a worker. Hence, these calculations need to be made on a two-dimension grid, where there is a range of values for both these dimensions.

In summary, when faced with a range of an investment strategy coupled with a guarantee structures, we propose that the analytical inputs which should inform the real world policy process is a quantification about four aspects: the price of the guarantee, the probability of falling into poverty, the median replacement rate, and the interquartile range of the replacement rate.

3 "Cassandra" - a software system for analysing guarantees

3.1 The model

We focus upon investment risk and the valuation of guarantees in the context of a simplified model.

Timeframe A worker embarks on pension contributions at time 0, and retires at age T.

Income We assume that his wage starts off at W_0, and grows at a fixed rate of w percent per year. This rate of growth reflects a combination of economic growth in the country, and the evolution of each worker upon the wage-experience profile prevalent in the country. We assume that there is no uncertainty about future wages.

Contribution rate We assume a fixed contribution rate c for life, whereby cW_t is accreted into the pension account in each time period.

Inflation All our analysis is conducted in real terms. This has the strong advantage that all numerical values seen in this paper, whether they pertain to contributions at age 20 or benefits at age 60, are interpretable in terms of price levels prevalent today.

Equity versus government bonds We assume a fraction λ of the portfolio is always invested in equities. That is, we assume that the portfolio at time t consists of λV_t equity and $(1-\lambda)V_t$ in debt. We assume that all equity investments are implemented through an index fund (Shah and Fernandes, 2001).

Asset market assumptions We assume that daily log returns on the equity index are drawn from $N(\mu, \sigma_M^2)$, and that government bonds risklessly yield a non-stochastic return r_f.

Annuity market assumptions We assume there is a market for annuities, where a worker of age T can convert a stock of assets A into a flow of LCU 1 per day until death. We assume complete annuitisation, i.e. on date T, the entire stock of pension wealth V_T is used to buy an annuity. We assume that consumption in old age is solely based on this annuity.

In the future, improvements in mortality (and fluctuations of interest rates) could yield different values for A. However, we assume that A is a constant. Alternatively, the value for A that is used in the calculations should reflect a forecast of the price of the annuity at pertinent future dates.

Poverty We assume that consumption below LCU $\bar{C}$ per day constitutes a state of poverty. In other words, an outcome $V_T/A < \bar{C}$ is synonymous with ending up in poverty in old age. Hence, we use the symbol $\bar{V}_T = \bar{C}A$.

The model works using iterations in time through the equations –

$$
\begin{aligned}
r_{Mt} &\sim N(\mu, \sigma_M^2) \\
V_{t+1} &= e^{r_{Mt}}\lambda V_t + (1+r_f)(1-\lambda)V_t + cW_t \\
W_{t+1} &= (1+w)W_t
\end{aligned}
$$

These iterations proceed, one day at a time, from time 0 till T. At date T, the stock of assets V_T is used in buying an annuity V_T/A giving a replacement rate $V_T/(AW_T)$.

Each guarantee structure corresponds to certain contingent payoffs. These can be valued using the general methods of the option pricing literature (Black and Scholes, 1973; Merton, 1973), which have been extensively used in research on pensions (Lachance and Mitchell, 2002; Pennacchi, 1998; Smetters, 2000).

These approaches to pricing critically rely on the assumption of complete markets, i.e. the assumption that financial markets offer trading opportunities for all states of nature. This assumption is clearly violated in the real world. Yet, pricing guarantees through this approach does constitute a useful benchmark for the price that would prevail in a certain idealised setting.

The key result of the option pricing literature is a concept called *risk neutral valuation*, which asserts that when markets are complete, the valuation of contingent claims can be done without reference to preference parameters. In the Black-Scholes world, the fair value of a contingent claim is the present value of payoffs under a risk-neutral random walk of underlying asset prices. The general approach to using simulation for measuring the value of the contingent claim is as follows:

1. Using the risk-free measure, we simulate sample paths of the underlying state variables over the relevant time horizon,
2. We estimate the discounted cashflows of a contingent claim on each sample path, as determined by the structure of the contingent claim in question, and
3. We average the discounted cashflows over sample paths.

Our software simulates a large number (typically, 5,000) trajectories through time, while assuming that there is no equity premium.[1] The NPV of the payout for each guarantee structure, in this 'risk neutral world', is the fair value of the guarantee.

Our analysis is a partial equilibrium analysis, in the sense that we do not consider the consequences upon equilibrium rates of return of a large volume put-protected equity investments by the pension system (Constantinides *et al.*, 2002). We return to this issue in Section 3.6.

3.2 Parameters about a country

The simulation requires seven numbers about a country of interest.

[1] It is important to distinguish between the 'real world' and the 'risk neutral world'. When we seek to understand issues like the probability of encountering poverty in old age, or the inter-quartile range of the replacement rate, we need to conduct simulations in the real world, where the magnitude of the equity premium matters. For the limited purpose of computing the value of contingent claims, we work in the risk-neutral world, where our assumption about the equity premium is irrelevant.

1. *Range of wages of interest.* As emphasised above, the simulation is conducted on a grid of plausible values for the initial wage per day. The low end of these wages should reflect the income of genuinely poor people, who are often particularly important in debates about pension reform. The upper end of these wages should reflect a typical civil servant wage.

   As a thumb rule, the left edge of the grid could be a value which is one-sixth of the per capita GDP of the country, and the right edge of the grid could be 50% bigger than the per capita GDP of the country.

   As an example, if a country has a per capita GDP of LCU 30,000 per year, a plausible grid of wages of interest could run from 20 to 200 LCU per day, in steps of 10 LCU.

2. *Average annual wage growth in real terms per year.* When the life of a pension system participant is being simulated, future wages are simulated to grow at a constant percentage rate. This growth reflects the sum of the rise in wages found through the wage-experience profile, and the outlook for per capita GDP growth in the country.

3. *Projection for the short-term riskless interest rate (r_f) in real terms.* This requires an estimate of what the short-end of the yield curve will be, expressed in real terms in percent per year. In developing countries, high values such as 2.5 percentage points are found.

4. *Projection for the equity premium.* This requires an estimate of what the returns on the equity index will be, over and above the short-term real riskless interest rate (r_f). Ideally, this should reflect the historical experience with a long time-series of 40 years or more for a well maintained stock market index. This is infeasible in developing countries. In this case, shorter data should be utilised to make a judgmental modification upon global thumb rules, which involve an equity premium of around 7 percentage points in the typical developing country.

5. *Projection for the daily volatility of the equity index.* Even if a well constructed stock market index is available for only a short time period, such as 10 years, it can be utilised to compute the standard deviation of daily returns on the index in a fairly accurate way.

6. *The price of an annuity at the age of retirement, which delivers a flow of LCU 1 per day until death.* This is the price of a "pure annuity", with no fund management services bundled in, sold by a traditional life insurance company.

7. *The threshold of poverty when understanding post-retirement income measured in LCU/day.* Many experts on poverty believe that USD

1 per day constitutes an interesting and objective definition of the poverty line. However, with CASSANDRA it is easy to use alternative thresholds of poverty.

3.3 Generic parameters

Apart from the parameters discussed above, which are specific to a country of interest, there is an additional set of parameters which are about the model and the simulation framework, and are not related to the country being analysed. They are:

1. The number of working days per year
2. The length of the working life in years - i.e. the gap between the year at which the worker enters the labour force and the year in which retirement takes place.
3. The contribution rate into the DC pension system. This is assumed to be a constant scalar through the entire life.
4. The number of simulations of the lifetime trajectory that are made for the purpose of evaluating a guarantee.

3.4 Software implementation issues

The calculation of a range of guarantee structures, on a two-dimensional grid where the equity fraction and the initial wage are varied, poses a considerable computational challenge. Naive implementations can easily run into weeks of computational time on conventional desktop computers.

Two key insights enable efficient implementation. The first aspect is to unroll the loop involved in the iterations from time 0 to time T, and re-express the task as a dot product. When there are non-stochastic contributions C_i, and if the initial wealth starts at 1, W_t evolves through equations such as:

$$\begin{aligned}
W_0 &= 1 \\
W_1 &= C_1 + W_0(1 + r_1) \\
W_2 &= C_2 + (1 + r_2)(C_1 + W_0(1 + r_1)) \\
W_3 &= C_3 + (1 + r_3)(C_2 + (1 + r_2)(C_1 + W_0(1 + r_1)))
\end{aligned}$$

The expression for W_3 can be rewritten as:

$$W_3 = C_3 + (1+r_3)C_2 + (1+r_3)(1+r_2)C_1 + (1+r_3)(1+r_2)(1+r_1)$$

which is a dot product between the vector $[1, (1+r_3), (1+r_3)(1+r_2), (1+r_3)(1+r_2)(1+r_1)]$ and $[C_3, C_2, C_1, 1]$.

This allows us to replace the computation of W_i in a loop by a dot product. This improves efficiency by eliminating loop overheads, and more importantly, enabling vectorised code which can use an efficient underlying Basic Linear Algebra System (BLAS) such as one produced by ATLAS or a hand-tuned BLAS supplied by the hardware vendor.

The second insight is to see that the generation of $N(0,1)$ random numbers is an extremely costly operation. Hence, we pre-generate a large number of $N(0,1)$ random vectors, and hold these in core for the purpose of all the simulations. When a random time-series of r_M is required, it is created using these prestored random variates through a calculation of $\mu + \sigma z$. Profiling studies reveal that computing a $N(0,1)$ vector is roughly five times costlier than computing $\mu + \sigma z$ using a pre-stored z.

These calculations have been implemented in a software package named CASSANDRA which is now available in source code form in the public domain at the URL `http://www.mayin.org/ajayshah/papers/cassandra` on the world wide web. This makes the results of this paper exactly reproducible, and enables the use of the suggested methods by other researchers. This software system is written in R, the statistical system which is also available in the public domain. Thus, both R and CASSANDRA are freely available.

3.5 Example: India

The numerical values that we use in this paper, which appear to be plausible in an Indian setting, are:

- T We assume $T = 40$, which may be roughly equivalent to a labour market career which runs from age 20 to age 60. This implies simulating over 10,000 days in one trajectory.
- w Wage growth is assumed to be 3% per annum.
- c The contribution rate is assumed to be 8.33%.
- r_f We assume that the real return on government bonds will prove to be 3% per year.
- r_M We assume the daily returns on the index has volatility 1.3% (per day). We assume that the equity premium will prove to be 7 percentage

points. These values are plausible estimates for the NSE-50 index (Shah and Thomas, 1998), which is the dominant index used with index derivatives and index funds in India.

A We assume that the market for annuities sells an annuity for a price of Rs.3,842 which yields a flow of Rs.1/day. This is drawn from the price of annuities sold by the Life Insurance Corporation (LIC), as of 2002.

$\bar{C}$ We assume that the threshold of poverty is $\bar{C} = 50$, which is roughly equivalent to consumption of a dollar a day. Conversely, the worker needs $\bar{V}_T = 192,100$ in order to avoid poverty in old age.

The results offered by CASSANDRA are in the form of a panel of figures which facilitate graphical visualisation. Each of these figures has the equity fraction and the initial wage as two axes, and a third parameter which is represented through colour and using contour lines.

Once equity investment is in the picture, the replacement rate in old age is a random variable. This is summarised in three ways. First, Figure 1 shows the probability that a pension system participant ends up in poverty in old age. This is like the probability of falling in the left tail, to the left of an absolute definition of poverty.

The remarkable feature of the probability of poverty is the interplay between initial income and the equity fraction. For the poor, investment in government bonds generates a certainty of poverty, but equity investment offers some hope of escaping poverty. Thus we get a 'plateau' where low W_0 and low λ generates a near-certainty of ending up in poverty. In contrast, for the rich, investment in government bonds offers a certainty of escaping poverty, while equity investment generates some possibility of suffering from it.

Figure 2 shows the median value of the replacement rate obtained by the participant and Figure 3 shows the interquartile range of the replacement rate. The median value of the replacement rate serves as a location estimator of the replacement rate, and the interquartile range of the replacement serves as a measure of dispersion, or the ex-ante uncertainty about consumption in old age as seen by the pension system participant. Higher fractions of equity investment generates a larger dispersion of outcomes, and altering W_0 does not influence this dispersion.

Finally, Figure 4 shows the price of the guarantee, expressed in logs to the base 10. Thus when a value of 4 is shown, this means 10^4 local currency units.

The guarantee of avoiding poverty has a simple relationship with income. The cost of guaranteeing that a low-income worker will avoid poverty in old age is higher than that seen for the rich. At the same time, the equity

Figure 1 The probability of falling into poverty in old age

This figure shows the probability of being in poverty in old age, under four alternative guarantee structures. The vertical bar on the right shows a mapping scheme from colours to probabilities. As an example, white is used for probabilities between 0.4 and 0.6, pink is used for near-zero probabilities and green is used for near-1 probabilities.
All four panes use the identical x and y axes, where the x axis consists of varying the equity fraction from 0 to 1, and the y axis consists of varying the initial wage from 20 to 200 local currency units.
The left hand bottom pane pertains to the situation where there is no guarantee. It shows a plateau of poverty for people with a low initial wage and low equity investment. As the equity fraction becomes higher, and as the initial wage becomes higher, the probability of poverty drops sharply.
Under structure G1, there is obviously no poverty, by construction. Under structure G2, where a replacement rate of 50% is guaranteed, there is no poverty except for the very poor, who end their labour market career with very low wages. Finally, a guarantee of an above-zero real rate of return has only a small effect upon the probability of poverty when compared with the state with no guarantee.

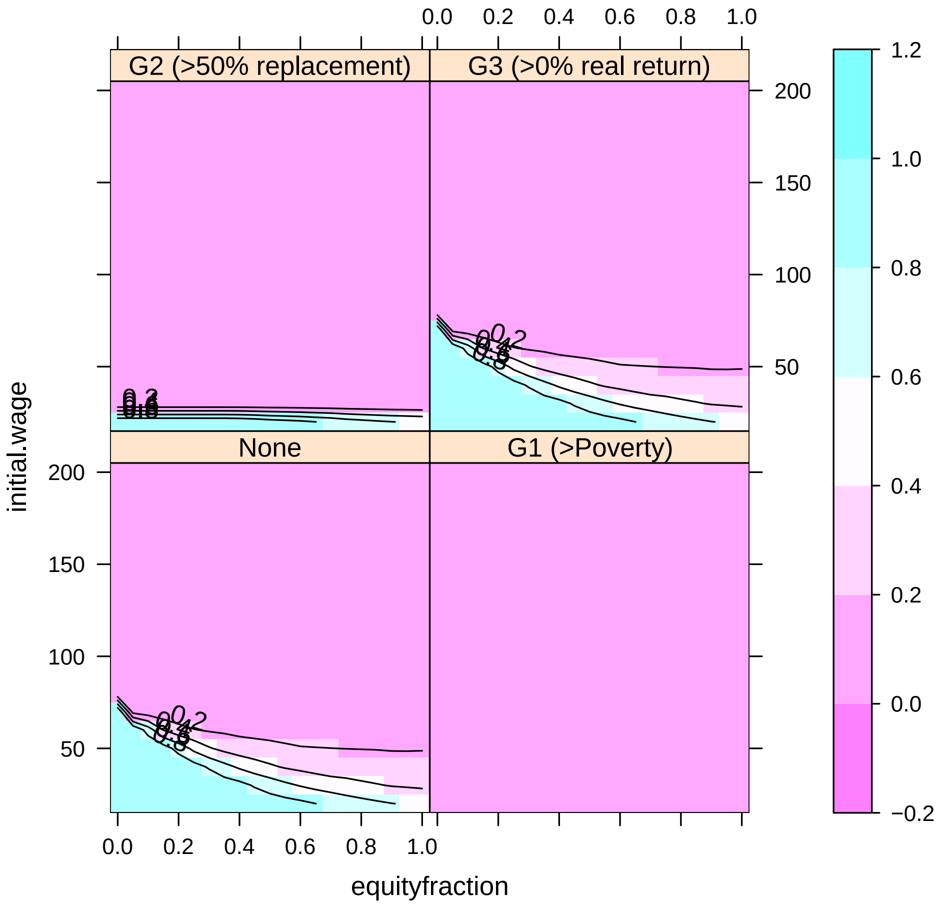

Figure 2 The median replacement rate

This figure shows the median replacement rate, where the median is computed across the distribution of potential outcomes - this serves as a location estimator. The vertical bar on the right shows a mapping scheme from colours to replacement rates. As an example, white is used for a replacement rate between 40% and 50%, pink is used for near-zero replacement rates and green is used for high replacement rates.
All four panes use the identical x and y axes, where the x axis consists of varying the equity fraction from 0 to 1, and the y axis consists of varying the initial wage from 20 to 200 local currency units.
The left hand bottom pane pertains to the situation where there is no guarantee. The contour lines show that the replacement rate is influenced only by the equity fraction and not by the initial wage. It goes from 20% for no equity investment to 60% for high equity investment.
Guarantee structure G1, which tops up above the poverty line, drives up the replacement rate for people with a low initial wage. Guarantee structure G2, which guarantees a 50% replacement rate, removes all the situations where the replacement rate might have dropped below 50%. Guarantee structure G3 does not substantially alter the unguaranteed outcome.

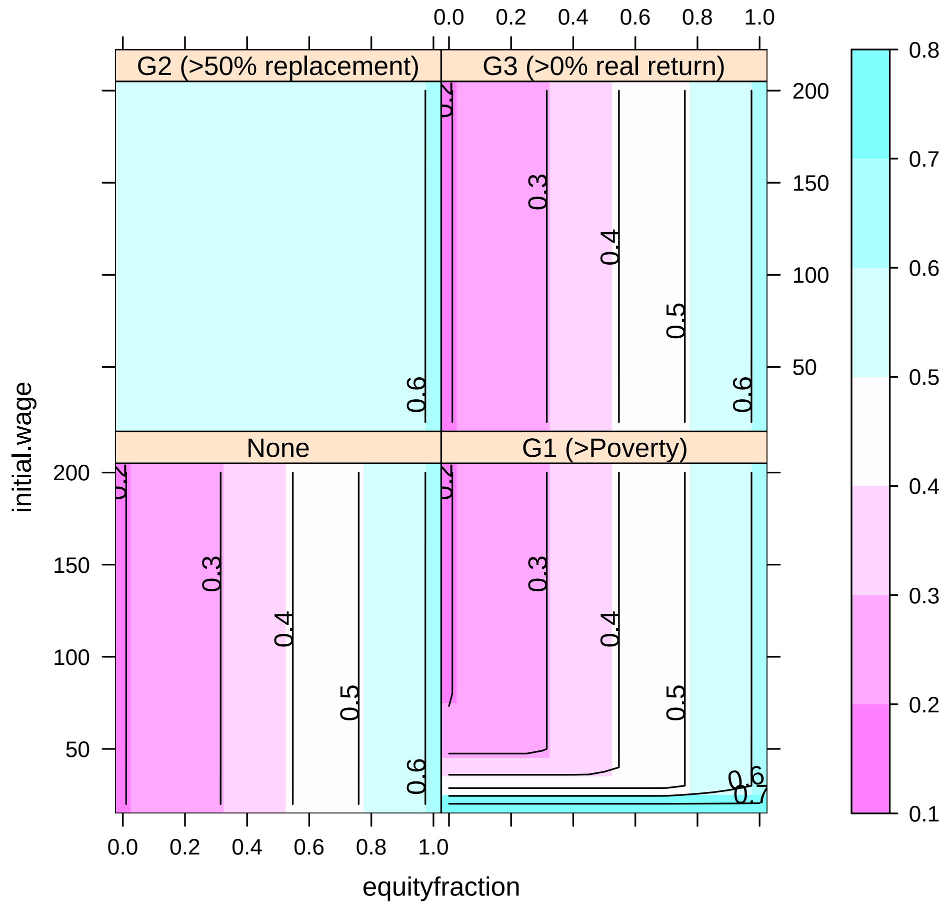

Figure 3 The interquartile range of the replacement rate

This figure shows the interquartile range of the replacement rate, computed across the distribution of potential outcomes - this serves as a measure of ex-ante investment risk as seen by the pension system participant. The vertical bar on the right shows a mapping scheme from colours to IQR. As an example, white is used for an IQR between 0.3 and 0.4, pink is used for near-zero risk and green is used for high levels of risk.
All four panes use the identical x and y axes, where the x axis consists of varying the equity fraction from 0 to 1, and the y axis consists of varying the initial wage from 20 to 200 local currency units.
The left hand bottom pane pertains to the situation where there is no guarantee. The contour lines show that risk is influenced only by the equity fraction and not by the initial wage. The interquartile range goes from values below 0.1 for equity investment below 30% and rises with higher levels of equity investment.
Guarantee structure G1, which tops up above the poverty line, eliminates investment risk for the poor, even if they have substantial equity investment. Guarantee structure G2, which ensures a 50% replacement rate, pushes all the contour lines to the right, delivering an IQR of 0.1 at a higher equity fraction of 60%. The G3 guarantee structure is similar to the unguaranteed case.

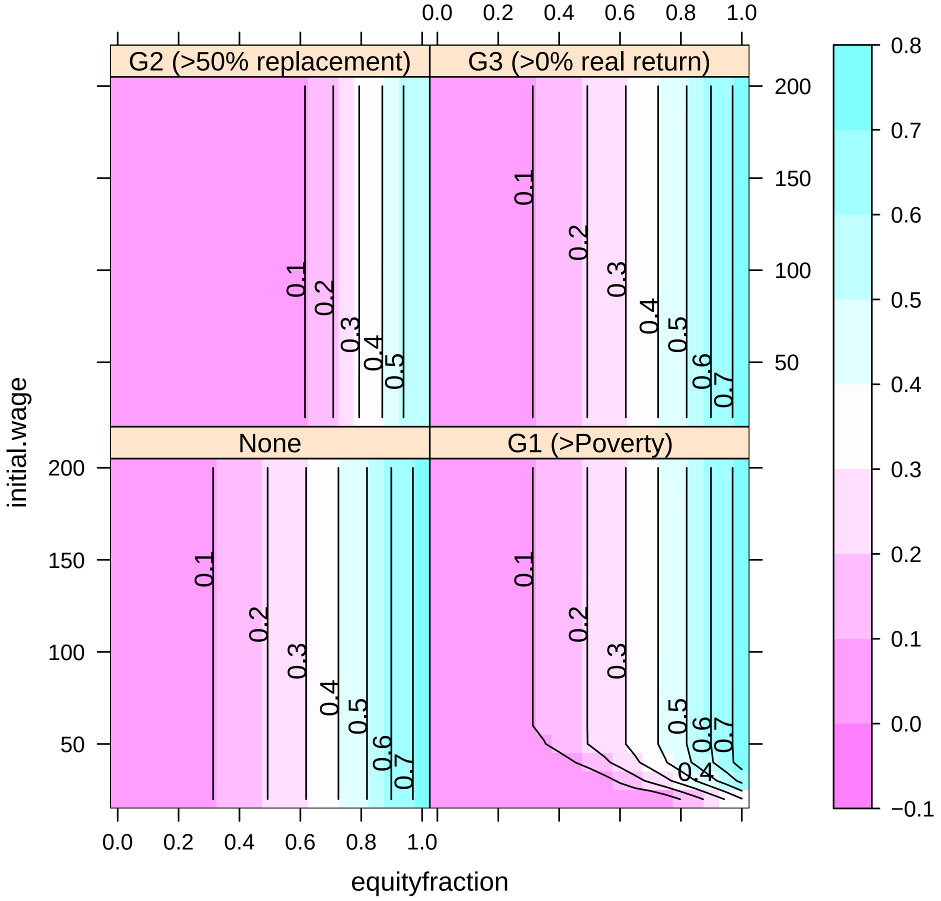

Figure 4 The price of the guarantee (shown in log to the base 10)

This figure shows the price of the guarantee. The previous figure can be viewed as measuring the benefits of guarantees through their impact on poverty, the median replacement rate and the uncertainty of the replacement rate. This figure shows the price paid to achieve these goals.

The vertical bar on the right shows a mapping scheme from colours to price. The unit used in log to the base 10 of the price expressed in local currency units, so a value of 5 here means a price of 10^5. As an example, white is used for a price between 100 and 1000 LCU, pink is used for a near-zero price and green depicts prices from 100,000 to a million. All four panes use the identical x and y axes, where the x axis consists of varying the equity fraction from 0 to 1, and the y axis consists of varying the initial wage from 20 to 200 local currency units.

The left hand bottom pane pertains to the situation where there is no guarantee, so the price is obviously always 0. Guarantee structure G1, which tops up above the poverty line is expensive for the poor and expensive when there is a large equity fraction. The price of this guarantee goes all the way to 10^4. The G2 guarantee structure is extremely expensive, with a cost of above 10^5 for people with an initial wage of above 75 LCU or so. The guarantee of an above-0 real return is worth a lot to high-income people with a high equity fraction: there is a right-hand corner with values of above 10^4.

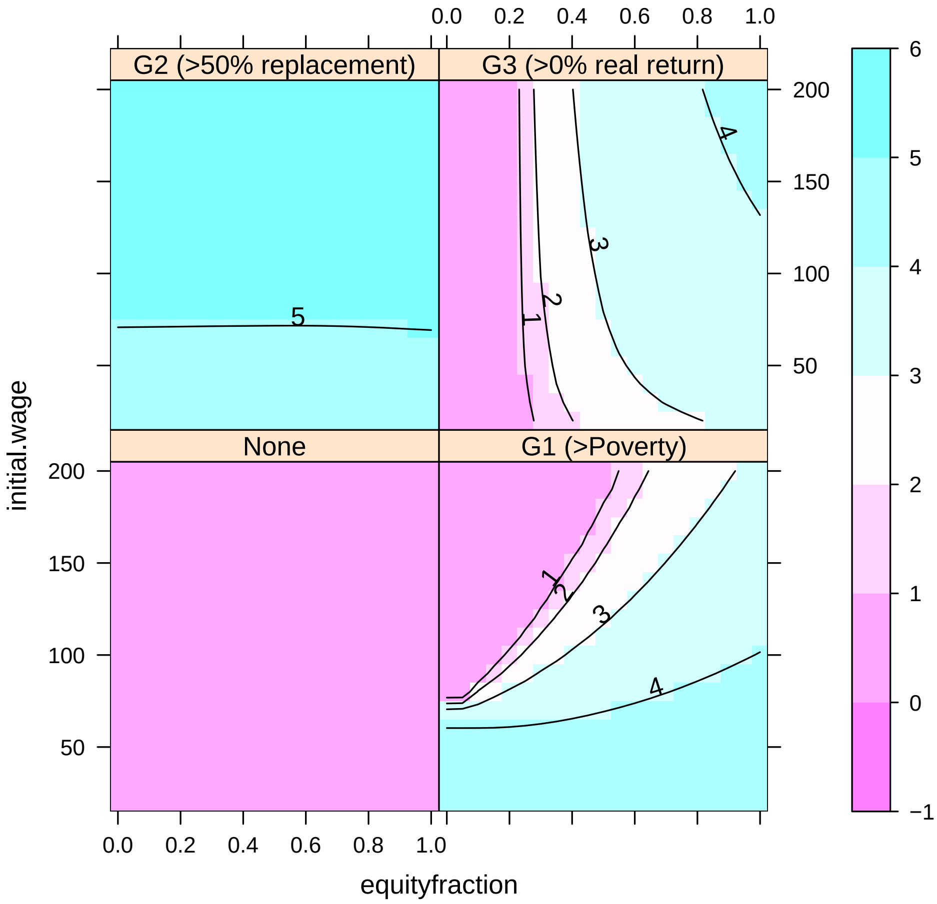

fraction affects this also. For poor people, a higher equity fraction reduces the cost of the guarantee. On the other hand, rich people get a downside protection when given this guarantee.

When it comes to a guarantee that the real returns will exceed 0%, the prices faced are substantially smaller than those seen for the guarantee that poverty will not occur. Conceptually, over these long investment horizons, it is extremely unlikely that a series of negative equity returns will be encountered, thus generating an overall negative real return. Hence, this guarantee is relatively unlikely to be exercised, and costs little. As expected, we see that the price of this guarantee is higher for the rich, and for those who invest more in equities.

The guarantee that the replacement rate will exceed 50% is an extremely expensive one. The price of this guarantee reflects the relationship between our assumptions for wage growth, contribution rate and the riskless rate of return. It is possible to obtain lower estimates for the price of this guarantee by assuming low rates of wage growth, high contribution rates, and high rates of return.

This structure leads to particularly high prices of the guarantee for the rich (who endup with a large value for W_T), and for those investing in equities (who keep the upside of the equity market, but pass on the downside risk to the provider of the guarantee).

In summary, capital protection seems to be the cheapest of the guarantees. The no–poverty guarantee) is expensive for the poor, and cheap for the rich (who are likely to escape poverty based on their own contributions). In contrast, the replacement rate guarantee is very expensive for the rich. All guarantees are costlier with high λ, since they involve options on a more volatile portfolio. Finally, if these guarantees are scaled up on the scale of an economy, they add up to truly enormous numbers.

3.6 Limitations

The strength of CASSANDRA lies in the fact that quantitative estimates about four important summary statistics are obtained for four alternative guarantee structures. However, this analysis has many limitations.

The most important limitation concerns the framework of risk-neutral valuation. The theoretical idea of risk-neutral valuation is applicable in a complete-markets world, where certain kinds of dynamic trading strategies can be frictionlessly executed. In developing countries, these assumptions are likely to be frequently violated. Hence, the price of the guarantee obtained through this analysis should only be viewed as that applicable in an idealised world.

The framework underlying CASSANDRA is one where the real returns on government bonds are deterministic. Investment risk only comes into the model through equities. This generates a bias against equities, insofar as developing countries have substantial volatility of inflation and interest rates.

The investment strategy that is analysed in CASSANDRA is where one decision – the fraction invested in equities – is made for life. In the real world, more sophisticated strategies can be adopted, which could involve higher equity exposure when young and a reduction in equity exposure later in the life cycle. While CASSANDRA can be modified to analyse such investment strategies, the existing style of computations are more revealing in explicitly showing the links between equity investment and the costs and benefits of guarantees.

One weakness of the approach of this paper lies in the fact that a set of parameters need to be established which will accurately reflect conditions over coming decades. This introduces model risk, for the extent to which the future will stray away from maintained parameters. This problem can be partly alleviated by varying all parameter values and checking for sensitivity.

The analysis of CASSANDRA is conducted in partial equilibrium. We assume that the data generating process for stock prices, and the fair value of option prices, is unaffected by the design of the pension system. However, quite possible that when a population-wide pension system enables large-scale purchases of the index spot and of index put options, prices on these markets will be distorted given the disproportionate size of pension assets when compared with the assets controlled by arbitrageuers.

The Black/Scholes model, and modern ideas on the pricing of derivatives, are founded on the no-arbitrage principle. The model implies certain links between the spot and the derivatives markets. If a high volume of buying of put options makes them 'too costly', this will throw up arbitrage opportunities, and private traders will exploit these arbitrage opportunities, so as to bring the spot price and volatility in line with the option prices.

Many developing countries have barriers against the operation of hedge funds, and have a framework of financial regulation which inhibits arbitrage activities of banks, insurance companies and pension funds. Assuming these problems will be addressed, and quasi-infinite global capital is in play engaged in derivatives arbitrage, we may expect that spot prices and option prices will be internally consistent, regardless of the scale of purchase of put options by customers of the pension system. Yet, this leaves questions about the extent to which large equity investments by the pension system can possibly modify the equity premium the volatility of the equity index - thus violating an assumption underlying CASSANDRA.

4 Policy issues of operationalising guarantees

4.1 The fiscal dimension

In debates on pensions, it is often assumed that the State would fund guarantees using general tax revenues. This approach suffers from several weaknesses:

Fiscal impact There are two aspects of the fiscal impact of guarantees.

First, if the State undertakes guarantees, but does not hedge its exposure using derivatives markets, then the stock of outstanding guarantees constitute off-balance sheet contingent liabilities for the State. This introduces non-transparency and risk in public finance.

Further, even if the State engages in sound risk management and reporting, where guarantees are implemented using private markets, our calculations suggest that giving free guarantees to workers is extremely expensive. For a pension system with 100 million participants, a guarantee worth \$1000 per person constitutes a transfer of \$100 billion from the exchequer to pension system participants.

Incentive implications The behaviour of pension participants would not be unchanged when a guarantee program is announced. As highlighted above, in a DC pension system, guarantees could generate actions which generate lower contributions and/or a higher fraction of assets invested in equities. For example, when there is a guarantee of a 50% replacement rate at retirement date, workers have an incentive to work for fewer years, and concentrate on only having a high terminal wage rate.[2]

The moral hazard induced by pension guarantees is also closely linked to population-wide anti-poverty programs. If there is a safety net, then individuals are more likely to engage in high risk bets, where they keep the upside, but can fall back upon State support in the event that low returns materialise. This suggests that the poor are particularly likely to exploit investment strategies with high equity investment and low expenditure on guarantees, if there is a credible assurance of a safety net from the State.

[2] In an extreme scenario, under guarantee structure 3, a worker who choose to be unemployed for life, but only works at a high wage for one month at age 60, would get an attractive pension for life. It may be assumed that in the real world, if such a guarantee structure is used, there would be safeguards against such behaviour. However, more modest versions of distorted behaviour would surely arise, such as efforts to work for fewer years and ensure a high terminal wage.

Political economy Guarantees raise concerns about fiscal prudence, and fiscal non-transparency. Many developing countries are placed in a situation with significant fiscal stress. In such an environment, if pension guarantees were a tool of public policy, there are concerns about imprudent guarantees being announced without corresponding risk management strategies and transparency about the fiscal implications.

Distributional implications Our results show that alternative guarantee structures have very different distributional consequences. If a government gives out a guarantee of a minimum replacement rate, this greatly favours the rich, while a guarantee of avoiding poverty favours the poor.

4.2 Implementation difficulties in the real world

If there was an attempt at creating guarantees purely using private markets, the major constraint faced in most developing countries would be a lack of adequate liquidity on the index options market. There is no country in the world where index options have significant liquidity for horizons exceeding a few years.[3]

There might be a role for the State to be a market maker in overcoming this 'missing markets' problem (Moss, 2002). This could take the form of a government selling put options to pension fund managers, or the government putting out two-way quotes on options markets.

4.3 Feasible policy alternatives

Despite these difficulties, there are mechanisms through which a DC pension system in a developing country can move towards access to some kinds of guarantees:

100% self-funding principle It would be prudent for the DC pension system to put decisions in the hands of each participant, where the participant makes a choice to determine his own guarantee structure (if any), pays a full price for it, and the guarantee is actually purchased from private financial markets. This would minimise the moral hazard, and fiscal difficulties, that flow from the introduction of guarantees.

[3]If a liquid index options market came about, there may be concerns about credit risk. If an economic agent X has sold an index put option with a ten-year maturity, this implies that the pension system customer is exposed to the credit risk of X for the next ten years. In practice, this is not a problem as long as the transaction enjoys the credit enhancement services of the futures clearing corporation.

Short dated guarantees This implies that the guarantees which can be offered in the forseeable future would have to be short-dated. If liquidity on the index options market builds up, it may be possible to obtain guarantees going 1-3 years out.

For many workers, an attractive lifetime investment strategy may be to hold 100% unhedged equities when young, and initiate protection using put options when nearing retirement. This strategy can be fully achieved using a sequence of short-dated guarantees.

Fostering the market for index put options The key factor which can impact upon the availability of guarantees from private markets would be rising sophistication and liquidity of the index option market. There are numerous policy impediments which impede liquidity of index options markets in developing countries. These include regulatory impediments against equity derivatives transactions by banks, insurance companies and pension funds, limitations of bank credit to arbitrageurs, limitations on banks engaging in arbitrage, a lack of short selling mechanisms on the spot market which are required to do arbitrage, etc.

International diversification The easiest way through which a DC pension system in a developing country can access an equity premium, while having access to liquid private option markets where guarantees can be obtained, is to exploit international diversification.

International diversification is a sound and useful idea in investment, regardless of the question of guarantees. An internationally diversified portfolio will (in general) be safer than any portfolio held in only one country. The lower volatility of internationally diversified portfolios implies that guarantees would be cheaper when applied to these.

The availability of liquid options markets on offshore stock indexes constitutes one additional reason why international diversification should play a major role in pension policies of developing countries.

5 Conclusion

Pension guarantees are financial derivatives. As with all derivatives, the guarantee is associated with a price, and we have the ability to *measure* this price in a complete-markets setting. This paper has emphasised that pension guarantees have some complex ramifications, particularly in a developing country context.

In this paper, we have proposed a systematic approach to thinking about the costs and benefits of alternative investment strategies and guarantees.

The costs are summarised through measures of investment risk – interquartile range of the replacement rate; probability of falling into poverty; cost of guarantee – and the benefits are summarised through the median replacement rate. We have proposed that these four numbers should be utilised as summary statistics in the public policy discourse about asset allocation and guarantees, with a two-dimensional grid over a range of values for the equity exposure and for initial wage.

This paper also serves as documentation for CASSANDRA: a freely available software system that enables the calculation and visualisation of these measures for any country. Some insights are offered on the complexities of calibrating parameters for a developing country, where data time-series are short and many markets are repressed. The results of these calculations can help improve the understanding amongst policy makers about the costs and benefits of alternative investment strategies and guarantee structures. They should help more informed decision making on the part of the State. If a State adopts such liabilities, this could be done with a full awareness of the off-balance sheet consequences which flow from a stated guarantee structure. When a full assessment of the public debt is made, such calculations would be useful in computing the implicit pension debt which should be taken into account.

Our calculations highlight the extent to which individuals who are given free guarantees have incentives to alter their contribution rates and equity exposure, in order to maximise the value of the guarantee that is given to them. These distortions could have major consequences for the structure of asset prices in the economy, and for the finances of the government.

References

Black F, Scholes M (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, **81**, 637–654. 7

Constantinides GM, Donaldson JB, Mehra R (2002). "Junior must pay: Pricing the implicit put in privatizing social security." *Technical report*, NBER. 8

Lachance ME, Mitchell OS (2002). "Understanding individual account guarantees." *Technical Report 9195*, NBER. 8

Merton RC (1973). "Theory of rational option pricing." *Bell Journal of Economics and Management Science*, **4**, 141–183. 7

Moss DA (2002). *When All Else Fails: Government As the Ultimate Risk Manager*. Harvard University Press, 1st edition. 20

Pennacchi G (1998). "Government guarantees on pension fund returns." *Technical report*, World Bank. 8

Shah A, Fernandes K (2001). "The relevance of index funds for pension investment in equities." In R Holzmann, J Stiglitz (eds.), "New Ideas About Old Age Security: Toward Sustainable Pension Systems in the 21st Century," World Bank. 7

Shah A, Thomas S (1998). "Market microstructure considerations in index construction." In "CBOT Research Symposium Proceedings," pp. 173–193. Chicago Board of Trade. 12

Siegel J (2002). *Stocks for the Long Run*. McGraw-Hill, 3 edition. 3

Sin Y (2002). "Minimum pension guarantees." *Technical report*, World Bank. 4

Smetters K (2000). "The design and cost of pension guarantees." *Technical report*, The Wharton School, University of Pennsylvania. 8

Walliser J (2002). "Retirement guarantees in mandatory DC systems." *Technical Report WP 2002-20*, Pension Research Council, University of Pennsylvania. 4