

Defined Contribution Pension Plans: Who Has Seen the Risk?

Peter A. Forsyth*

Kenneth R. Vetzal†

Draft version: February 21, 2019

Abstract

The trend towards eliminating defined benefit (DB) pension plans in favour of defined contribution (DC) plans implies that increasing numbers of pension plan participants will bear the risk that final realized portfolio values may be insufficient to fund desired retirement cash flows. We compare the outcomes of various asset allocation strategies for a typical DC plan investor. The strategies considered include constant proportion, linear glide path, and optimal dynamic (multi-period) time consistent quadratic shortfall approaches. The last of these is based on a double exponential jump diffusion model. We determine the parameters of the model using monthly US data over a 90 year sample period. We carry out tests in a synthetic market which is based on the same jump diffusion model and also using bootstrap resampling of historical data. The probability that portfolio values at retirement will be insufficient to provide adequate retirement incomes is relatively high, unless DC investors adopt optimal allocation strategies and raise typical contribution rates. This suggests there is a looming crisis in DC plans, which requires educating DC plan holders in terms of realistic expectations, required contributions, and optimal asset allocation strategies.

JEL Classification: G11

Keywords: defined contribution plan, probability of shortfall, mean-variance, dynamic asset allocation, resampled backtests

1 Introduction

Defined benefit (DB) pension plans can be high risk liabilities for both private and public sector organizations. To de-risk balance sheets, many employers have been gradually converting to defined contribution (DC) plans. This trend has been particularly strong in Australia and the US: as of 2017, 87% of pension assets in Australia were in DC plans, while the corresponding figure for the US was 60%. In contrast, in countries such as the UK, Canada, and Japan pension plan assets were still predominantly in DB plans.¹ Of course, considering just the percentage of assets in DC plans at a point in time can hide the underlying trend. For example, about 95% of Canadian pension assets were in DB plans as of 2017. However, Gougeon (2009) points out that the number of DC

*David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, paforsyt@uwaterloo.ca, +1 519 888 4567 ext. 34415. This work was supported by the Natural Sciences and Engineering Research Council of Canada.

†School of Accounting and Finance, University of Waterloo, Waterloo ON, Canada N2L 3G1, kvetzal@uwaterloo.ca, +1 519 888 4567 ext. 36518.

¹Statistics from the Thinking Ahead Institute's "Global Pension Assets Study 2018". Available online at www.thinkingaheadinstitute.org/-/media/Pdf/TAI/Research-Ideas/GPAS-2018.pdf.

plan participants in Canada almost doubled between 1991 and 2006, while the number of DB plan participants shrank by about 4% over the same period.

Under a DC plan, the employee and employer contribute to a retirement savings account, which is usually tax-advantaged. In some cases, the employee is able to select from a list of investments (e.g. stock and bond index funds) which have been approved by the employer. However, in terms of the important decisions about asset allocation, the employee is typically left up to her own devices.

Although DC plans are desirable from the employers' point of view, the retirement savings risk has simply been transferred to those who are perhaps least able to manage it. Many studies have confirmed that the average retail investor is very poor at managing investment portfolios (e.g. Barber and Odean, 2013). A typical DC plan investor will accumulate savings for thirty or more years, followed by a decumulation phase of perhaps twenty years. As a result, inefficiencies in asset allocation, even if small in annual terms, can have very large cumulative effects over a typical long term retirement savings investment horizon.

Target date funds (TDFs) have become a popular means of providing an *autopilot* asset allocation strategy. The prototypical TDF uses a high allocation to equities during the early years of DC savings. The stock allocation is decreased, and the portfolio becomes more weighted towards bonds, as the retirement date nears. The transition from high to low equity allocations is pre-determined by means of a *glide path* strategy. In the US, TDFs are *Qualified Default Investment Alternatives*, implying that the assets of employees enrolled in an employer-managed DC plan can be invested in a TDF by default.² Close to \$900 billion was invested in US TDFs as of the end of 2016 (Holt, 2017).

However, several studies have reported that TDFs do not appear to outperform simpler constant proportion strategies (Basu et al., 2011; Arnott et al., 2013; Esch and Michaud, 2014; Estrada, 2014; Forsyth and Vetzel, 2019). Clearly, constant proportion strategies which have a significant allocation to equities could result in a wide range of terminal wealth at the end of the accumulation period. The fact that TDFs are not superior to such strategies indicates that TDFs are not a panacea for DC plan investors.

The objective of this article is to describe realistic outcomes for a typical DC plan investor, who contributes a fixed (real) amount per year over a lengthy (30 year) accumulation period. We assume that the investor has two possible assets in the DC account: a low risk bond index fund and an equity index fund. We model the real (inflation-adjusted) stock index as following a double exponential jump diffusion model (Kou and Wang, 2004), which incorporates the risk of sudden market crashes.³ The parameters of the jump diffusion model are estimated by fitting to nine decades of US market data.

We consider three different asset allocation strategies: (i) constant proportion; (ii) a linear glide path; and (iii) an optimal dynamic (multi-period) quadratic shortfall (QS) strategy. The optimal QS strategy is adaptive, i.e. the asset allocation is a function of the current state of the investment portfolio. This strategy is computed by numerically solving a Hamilton-Jacobi-Bellman (HJB) equation (Dang and Forsyth, 2014). It is interesting to note that the QS strategy, which is naturally time consistent, is equivalent to the pre-commitment multi-period mean-variance (MV) strategy, in the sense that both objective functions result the same optimal controls. This can be proven based on the embedding technique (Li and Ng, 2000; Zhou and Li, 2000). Our numerical approach allows us to impose realistic no-leverage constraints. This avoids the impractical situation of a highly leveraged portfolio, which often results for unconstrained pre-commitment MV optimal

²In other words, the employee must explicitly decide to opt out of a TDF if she desires a different asset allocation strategy.

³An obvious extension would be a GARCH/stochastic volatility model, but Ma and Forsyth (2016) document that mean-reverting stochastic volatility effects are negligible for long term investors.

strategies (Lioui and Poncet, 2016)⁴.

As a benchmark case, we consider a constant proportion strategy involving two assets: a risk-free bond index and a risky stock index. The portfolio is annually rebalanced. The weights on the two assets are chosen to achieve a target expected level of real wealth at the end of the savings period. We then compare this strategy with a glide path and an optimal QS strategy, both of which are constrained to have the same expected value of final wealth as the constant proportion strategy. The glide path and constant proportion asset allocation strategies can be considered to be typical of many DC plan holders. The optimal QS strategy is effectively a *best case* scenario, in terms of reducing the probability of shortfall (over a wide range of the terminal wealth distribution) compared to glide path or constant proportion strategies (Forsyth and Vetzal, 2017, 2019).

We compare these strategies in a *synthetic* market, in which the equity price process is assumed to follow a jump diffusion with constant parameters (i.e. the average historical parameters). We also carry out tests using bootstrap resampling with actual historical data (Cogneau and Zakalmouline, 2013; Dichtl et al., 2016).

We make several assumptions which allow us to convert a desired final salary replacement ratio into a target real portfolio value after a 30 year contribution period. Our base case example assumes a target real final salary replacement ratio of 50%, a combined annual employee-employer contribution of 20% of real salary, and investments in short-term risk-free government bonds and a value-weighted equity index.

Our two main findings are as follows:

1. With typical glide path or constant proportion strategies, there is an unacceptably large probability of shortfall in terms of meeting the target final wealth goal.
2. Even with an optimal dynamic QS asset allocation strategy, there is still a fairly high probability of shortfall. This shortfall probability can be reduced to what we view as a reasonable level by increasing the total contribution rate or reducing the replacement ratio, compared to the base case. Another possibility is to substitute an equal-weighted equity index for the value-weighted index, but this may not work in practice due to higher costs associated with equal-weighted indexes, which are not recognized in our model.

Our results suggest that most DC plan holders are unlikely to be able to maintain their desired standard of living upon retirement. In our opinion, the true risks of DC plans have not been communicated in realistic terms to plan participants. It seems that this situation can be improved only by simultaneously lowering expectations, increasing savings rates, and improving asset allocation strategies.

A possible criticism of our findings is based on the fact that our computations use the past 90 years of US market data. Some would argue that real market returns going forward will be lower than historically observed values. However, this would only lead to worse results, so it could be argued that our conclusions should be even more pessimistic.

2 Formulation

We focus exclusively on a simple context with just two assets available in the financial market, namely a risky asset and a risk-free asset. In practice, the risky asset would be a broad market index fund. An investor saves for retirement at time T . The amounts that this investor's portfolio contains of the risky and risk-free assets at time t are denoted by S_t and B_t respectively. The

⁴Recall the time consistent QS strategy has the same controls as the pre-commitment MV strategy

investor's total wealth from the portfolio (i.e. the total value of the portfolio) at t is $W_t = S_t + B_t$. The fraction of total wealth invested in the risky asset is $p_t = S_t/W_t$. The investment period runs from the inception time $t = 0$ to the horizon date $t = T$. There is a set of $M + 1$ pre-determined *action times* denoted by \mathcal{T} ,

$$\mathcal{T} \equiv \{t_0 = 0 \leq t_1 \leq \dots \leq t_M = T\}. \quad (2.1)$$

At the horizon date $t_M = T$, the portfolio is liquidated. At each action time $t_i \in \mathcal{T}_1 = \mathcal{T} \setminus \{t_M\}$ (i.e. each action time prior to T), (i) an amount of cash q_i is contributed to the portfolio and then (ii) the portfolio is rebalanced.⁵

Let the instant before action time t_i be $t_i^- = t_i - \epsilon$, where $\epsilon \rightarrow 0^+$. Similarly, the instant after t_i is denoted by $t_i^+ = t_i + \epsilon$. To simplify notation, let $S_i^+ = S_{t_i^+}$, $S_i^- = S_{t_i^-}$, $B_i^+ = B_{t_i^+}$, $B_i^- = B_{t_i^-}$, $W_i^+ = W_{t_i^+}$, and $W_i^- = W_{t_i^-}$. Similarly, let $p_{t_i^+} = p_i$.

Between action times (i.e. $t \notin \mathcal{T}$), the value of the investor's portfolio will fluctuate in accordance with changes in the unit prices of the two assets. We assume a constant risk-free rate r , so that the evolution of the amount invested in the risk-free asset is

$$dB_t = rB_t dt, \quad t \notin \mathcal{T}. \quad (2.2)$$

The dynamics of the changes in the amount invested in the risky asset between action times are given by the jump diffusion process

$$\frac{dS_t}{S_t^-} = (\mu - \lambda\kappa) dt + \sigma dZ + d\left(\sum_{i=1}^{\pi_t} (\xi_i - 1)\right), \quad t \notin \mathcal{T}, \quad (2.3)$$

where μ is the (uncompensated) drift rate, σ is the volatility, dZ is the increment of a Wiener process, π_t is a Poisson process with intensity λ , and ξ denotes the random jump multiplier. When a jump occurs, $S_t = \xi S_{t-}$, and $\kappa = E[\xi - 1]$ where $E[\cdot]$ is the expectation operator. We assume that ξ_i are i.i.d. positive random variables characterized by a double exponential distribution (Kou and Wang, 2004). Given that a jump occurs, p_{up} is the probability of an upward jump and $(1 - p_{up})$ is the probability of a downward jump. The density function $f(y = \log \xi)$ is then

$$f(y) = p_{up}\eta_1 e^{-\eta_1 y} \mathbf{1}_{y \geq 0} + (1 - p_{up})\eta_2 e^{\eta_2 y} \mathbf{1}_{y < 0}. \quad (2.4)$$

We do not permit short sales of the risky asset, and we impose an upper bound L_{\max} on the leverage ratio that the investor is allowed to use. This means that $0 \leq p_i \leq L_{\max}$ for all action times $t_i \in \mathcal{T}_1$. Generally, with a DC account it is reasonable to specify $L_{\max} = 1$, ruling out the use of any leverage. Since the value of the risky asset follows the jump diffusion (2.3), if we allow leverage by setting $L_{\max} > 1$ the investor can become insolvent.⁶ We add the further constraint that if the investor becomes insolvent at any time, then trading stops and all positions in the risky asset are liquidated.⁷ In insolvency, debt accumulates until it is (possibly) eliminated by cash contributions. We emphasize that insolvency can only occur if leverage is allowed, i.e. $L_{\max} > 1$.

⁵As discussed below, in the case of an optimal QS strategy the investor may also withdraw cash from the portfolio at an action time.

⁶Since the investor rebalances her portfolio discretely, insolvency could also occur if $L_{\max} > 1$ in the special case of the model where jumps are ruled out ($\lambda = 0$), i.e. the value of the risky asset follows geometric Brownian motion.

⁷More precisely, suppose that insolvency occurs at time t , i.e. $S_t + B_t < 0$. Letting t^+ be the instant after t , then $B_{t^+} = S_t + B_t$ and $S_{t^+} = 0$.

2.1 Deterministic Glide Paths

TDFs generally use a deterministic glide path, where the asset allocation depends only on time. In our case, this would imply $p_i = p(t)$. One case is a *linear glide path*, with

$$p_i = p_{\max} + \frac{t_i \times (p_{\min} - p_{\max})}{T} \quad (2.5)$$

where p_{\max} and p_{\min} are parameters. Note also that a constant proportion strategy can be viewed as a deterministic glide path with $p = \text{const.}$ for all action times $t_i \in \mathcal{T}_1$.

Between action times, the amounts the investor has in the risk-free and risky assets follow the processes (2.2) and (2.3) respectively. Recalling that q_i is a cash contribution, at action times prior to the horizon date (i.e. $t_i \in \mathcal{T}_1$), we have

$$\begin{aligned} W_i^+ &= S_i^- + B_i^- + q_i \\ S_i^+ &= p_i W_i^+ \\ B_i^+ &= (1 - p_i) W_i^+. \end{aligned} \quad (2.6)$$

In the case of deterministic glide paths, closed form recursive expressions for the mean and variance of terminal wealth W_T are developed in Forsyth and Vetzal (2019). The cumulative distribution function (CDF) for W_T is computed using a Monte Carlo method. In our numerical tests below, we compare all strategies by fixing expected terminal wealth. Since we have closed form expressions for the mean, we determine the glide path parameters using a Newton iteration in order to enforce this condition.⁸

2.2 Adaptive Strategies

In contrast to deterministic strategies where the asset allocation depends only on time, adaptive strategies allow the asset allocation to depend on the prevailing state of the investment portfolio. Since we search for the optimal controls over all portfolios with the same wealth after cash injection (W_i^+), this means that $p_i = p_i(W_i^+, t_i^+)$. With an adaptive strategy, it can be optimal to withdraw cash from the portfolio (Cui et al., 2014; Dang and Forsyth, 2016). We denote this optimal cash withdrawal as $c_i \equiv c(W_i^- + q_i, t_i)$. Since we only allow cash withdrawals, $c_i \geq 0$. The control at action time t_i now consists of the pair (p_i, c_i) , i.e. after withdrawing c_i from the portfolio, rebalance so that the fraction invested in the risky asset is p_i .

For explanatory purposes, let us consider first consider a dynamic (multi-period) MV criterion

⁸For example, we can exogenously specify p_{\min} and find the value of p_{\max} which generates the desired expected terminal wealth via Newton iteration. Alternatively, we can exogenously set p_{\max} and numerically find the appropriate value of p_{\min} .

with a specified desired value of $E[W_T] = W_d$. The problem to be solved can be stated as

$$\begin{aligned} \min_{\{(p_0, c_0), \dots, (p_{M-1}, c_{M-1})\}} \quad & \text{Var}[W_T] = E[W_T^2] - W_d^2 \\ \text{subject to} \quad & \begin{cases} \text{At horizon } T: E[W_T = S_T + B_T] = W_d \\ \text{Between action times } t \notin \mathcal{T}: (B_t, S_t) \text{ follow processes (2.2), (2.3)} \\ \text{At action times } t \in \mathcal{T}_1: \\ \quad W_i^+ = S_i^- + B_i^- + q_i - c_i \\ \quad S_i^+ = p_i W_i^+, \quad B_i^+ = W_i^+ - S_i^+ \\ \quad p_i = p_i(W_i^+, t), \quad 0 \leq p_i \leq L_{\max} \\ \quad c_i = c_i(W_i^- + q_i, t) \quad c_i \geq 0 \end{cases} \end{aligned} \quad (2.7)$$

A criticism of the pre-commitment MV problem (2.7) is that it is *time inconsistent*. In other words, the investor has an incentive to deviate from the strategy computed at time zero (Basak and Chabakauri, 2010). However, in order to solve problem (2.7), we can use the embedding technique (Zhou and Li, 2000; Li and Ng, 2000). Consider a control set

$$P = \{(p_i(W_i^+, t_i^+), c_i(W_i^- + q_i, t_i)), i = 0, \dots, M-1\}.$$

Informally, if P^* is an optimal control for problem (2.7), then there exists a W^* such that P^* is also the optimal control for the following problem

$$\begin{aligned} \min_{\{(p_0, c_0), \dots, (p_{M-1}, c_{M-1})\}} \quad & E[(W^* - W_T)^2] \\ \text{subject to} \quad & \begin{cases} \text{At horizon } T: E[W_T = S_T + B_T] = W_d \\ \text{Between action times } t \notin \mathcal{T}: (B_t, S_t) \text{ follow processes (2.2), (2.3)} \\ \text{At action times } t \in \mathcal{T}_1: \\ \quad W_i^+ = S_i^- + B_i^- + q_i - c_i \\ \quad S_i^+ = p_i W_i^+, \quad B_i^+ = W_i^+ - S_i^+ \\ \quad p_i = p_i(W_i^+, t), \quad 0 \leq p_i \leq L_{\max} \\ \quad c_i = c_i(W_i^- + q_i, t) \quad c_i \geq 0 \end{cases} \end{aligned} \quad (2.8)$$

Problem (2.8) can be solved using dynamic programming methods.⁹

As noted above, it is optimal to withdraw cash from the portfolio under some conditions (Cui et al., 2012; Dang and Forsyth, 2016). Let

$$Q_\ell = \sum_{j=\ell+1}^{j=M-1} e^{-r(t_j - t_\ell)} q_j \quad (2.9)$$

be the discounted planned future contributions to the DC account at time t_ℓ . If

$$(W_i^- + q_i) > W^* e^{-r(T-t_i)} - Q_i, \quad (2.10)$$

then the optimal strategy is to

⁹If problem (2.7) is not convex, there may be solutions to problem (2.8) that are not solutions to problem (2.7). However, these spurious solutions can be eliminated in a straightforward way (Tse et al., 2014; Dang et al., 2016).

- (i) withdraw cash $c_i = W_i^- + q_i - (W^*e^{-r(T-t_i)} - Q_i)$ from the portfolio; and
- (ii) invest the remainder $(W^*e^{-r(T-t_i)} - Q_i)$ in the risk-free asset.

This is optimal because in this case $E[(W^* - W_T)^2] = 0$, which is the minimum of problem (2.8).

We refer to any cash withdrawn from the portfolio as *surplus cash* in the following. For the sake of discussion, we will assume that surplus cash is invested in the risk-free asset, but does not contribute to the calculation of the mean and variance of terminal wealth.

This target-based approach of problem (2.8) provides a reasonable objective on its own (Vigna, 2014; Menoncin and Vigna, 2017; Vigna, 2017). Solving (2.8) minimizes quadratic shortfall (QS) with respect to W^* , so we will refer to the resulting strategy as the *QS optimal* strategy below. However, this becomes even more compelling when we recall that the solution is also pre-commitment MV efficient. The solution then simultaneously minimizes *two* risk measures: variance around the desired $E[W_T]$ and QS with respect to W^* , as seen at time zero.

We emphasize that the fact that the pre-commitment MV policy is time inconsistent is irrelevant since we take the point of view that we are seeking the QS optimal control, from problem (2.8). Since the QS problem (2.8) can be solved using dynamic programming, the controls are trivially time consistent. The fact that the QS problem gives rise to time consistent controls, whereas the MV problem (2.7) is time inconsistent, is due to the fact that we fix W^* for the QS problem, for all time. At time zero, the MV problem controls and the QS problem controls are the same for W^* computed at time zero. At later times, this correspondence holds only if we allow W^* to change as a function of time. However, using a fixed W^* is intuitively reasonable for DC pension plan saving (Vigna, 2014; Menoncin and Vigna, 2017).

We note that there are techniques for forcing a time consistent constraint for the MV problem (2.7) (Bjork and Murgoci, 2010; Wang and Forsyth, 2011; Bjork and Murgoci, 2014; Bjork et al., 2014). However, we prefer the target based QS approach since it is relatively easy to communicate to end user investors (Vigna, 2014; Menoncin and Vigna, 2017). In addition, forcing the time consistent constraint can have result in non-intuitive strategies with strange features (Wang and Forsyth, 2011; Bensoussan et al., 2019).

We formulate problem (2.8) as the solution of a nonlinear Hamilton-Jacobi-Bellman (HJB) partial integro differential equation. See Dang and Forsyth (2014) for details concerning the numerical solution. Given an arbitrary value of W^* , we can solve problem (2.8) for the optimal control, which we denote by $P^*(W^*)$. Given the optimal control, cumulative distribution functions are easily found using Monte Carlo simulation. However, we seek the solution to problem (2.7), which is expressed in terms of a specified expected value $E[W_T] = W_d$. We determine the value of W^* for problem (2.8) which satisfies the constraint $E[W_T] = W_d$. We enforce this by a Newton iteration, whereby each function evaluation requires a solution of an HJB equation.

3 Data and Parameter Estimates

The underlying stochastic model outlined above in equations (2.2), (2.3), and (2.4) involved a constant risk-free rate r for the bond component and a double exponential jump-diffusion for the equity component. Estimation of the parameters of these equations follows the methods described in Forsyth and Vetzal (2019). These procedures are summarized briefly here for convenience. Readers interested in additional details are referred to Forsyth and Vetzal (2019).

We use monthly US data obtained from the Center for Research in Security Prices (CRSP) for

the period 1926:1 through 2015:12.¹⁰ Our base case uses the CRSP value-weighted total return index (which includes all distributions for all domestic equities trading on major US exchanges), along with the CRSP 3-month Treasury bill (T-bill) index. The original data are in nominal terms, but we convert them to real terms using the US CPI, also obtained from CRSP. We use real indexes since investors with long-term savings objectives such as funding retirement should concentrate on real (not nominal) wealth goals. For some tests, we use alternative underlying assets: the CRSP equal-weighted total return index (which invests the same amount in each component security, rather than weighting by market capitalization) and a 10-year US Treasury bond (T-bond) index.

Figure 3.1 provides plots of monthly real returns for the 3-month T-bill, 10-year T-bond, value-weighted total return, and equal-weighted total return indexes. For comparability, all four indexes are plotted with the same vertical axis scale. As expected, the two equity indexes exhibit much higher volatility, with occasional months having returns of large magnitude. This provides a measure of support to our modelling assumptions which assume a constant interest rate but a jump diffusion specification for the equity index. It is also interesting to observe that there was extremely high equity market volatility during the 1930s. By contrast, volatility during the period following the financial crisis that began in 2007 was comparatively mild.

Figure 3.2 graphs cumulative real returns for the four investment indexes. The vertical axis uses a logarithmic scale. This enhances visibility over time, as otherwise the dramatic growth in the equity indexes over the latter part of the sample obscures the behaviour of those indexes during the earlier periods and renders the behaviour of the two Treasury indexes all but invisible. All four indexes begin at a value of 100 at the start of 1926. The equal-weighted index ends up with the highest value. The historical outperformance of equal-weighting has been attributed to such portfolios having higher exposure to value, size, and market factors (Plyakha et al., 2014). It is also interesting to observe that the 10-year T-bond index had higher cumulative returns than the 3-month T-bill index, but that was entirely due to the post-1980 period: prior to then, the longer maturity T-bond index offered no cumulative advantage over the T-bill index.

Table 3.1 presents parameter estimates. A threshold method (Cont and Mancini, 2011) was used for the jump diffusion model. These parameter estimates were originally provided in Forsyth and Vetzal (2019), and are reproduced here for convenience. The estimates using the value-weighted equity market index imply an expected real annual return of close to 9%, about 3% lower than the corresponding value for the equal-weighted index. Of course, the price to be paid for this difference is higher risk. The equal-weighted index shows higher diffusive volatility (σ). Since jumps are expected to occur on average every $1/\lambda$ years, the equal-weighted index tends to have jumps a bit more often. Conditional on a jump occurring, it is much more likely to be a downward jump in each case. Average jump magnitudes are $1/\eta_1$ for upward jumps and $1/\eta_2$ for downward jumps, and these are both larger for the equal-weighted index. A similar comment applies to the standard deviation of the jump size since this is equal to the mean for the exponential distribution. Turning to the bond market indexes, Table 3.1 shows that the long run average real annual return for the 10-year T-bond index was just over 2%, while that for the shorter maturity index was around 80 basis points. Of course, these higher returns are accompanied by higher volatility, as indicated by the top two plots of Figure 3.1.

¹⁰More precisely, our calculations are based on data from Historical Indexes, ©2015 Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business. Wharton Research Data Services (WRDS) was used in preparing this work. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.




FIGURE 3.1: *Monthly real returns for US investment indexes, 1926-2015.*

Equity Market Index	μ	σ	λ	p_{up}	η_1	η_2
Value-weighted	0.0889	0.1477	0.3222	0.2759	4.4273	5.2613
Equal-weighted	0.1183	0.1663	0.4000	0.3333	3.6912	4.5409
Bond Market Index	Average Return					
3-month T-bill	.00827					
10-year T-bond	.02160					

TABLE 3.1: *Annualized parameter estimates based on real monthly data from 1926:1 to 2015:12. These values originally appeared in Forsyth and Vetzal (2019) and are reproduced here for convenience. Parameters for the equity market indexes were estimated using the threshold technique of Cont and Mancini (2011). The average returns for the bond indexes were calculated as $\log[B(T)/B(0)]/T$, where $B(t)$ denotes the index level at the time t .*




FIGURE 3.2: *Cumulative real returns for US investment indexes, 1926-2015.*

4 Base Case Scenario

We consider an example where a DC plan member wants to generate retirement income (in real terms) of a specified fraction R of final salary. Studies have shown that earnings for a typical employee increase rapidly until the age of 35, then increase slowly thereafter, until a few years before retirement, and then decrease as fewer hours are worked in the transition to retirement (Cocco et al., 2005; Ruppert and Zanella, 2015). Given an initial real annual salary I_0 , we assume that real annual income in year t before retirement is given by $I_t = e^{\mu_I t} I_0$. In other words, salary is expected to grow in real terms at a constant annual rate μ_I . Upon retirement at time T , real annual salary is $e^{\mu_I T} I_0$. Our assumptions would be relevant to a 35 year old employee with stable employment, who intends to work full-time until the age of 65.

To determine the amount of real wealth required to fund this replacement income during retirement, we use the well-known 4% rule of Bengen (1994). Bengen examined historical data to determine the maximum real withdrawal rate that a retiree could safely use without exhausting her assets over a 35 year period. Bengen assumed that accumulated pension wealth was invested in a portfolio having half invested in stocks and half invested in intermediate-term US Treasury securities, and concluded that a 4% withdrawal rate (escalated by the rate of inflation) was quite safe.

Dang et al. (2017) recently revisited this rule. The problem was posed somewhat differently: the idea was to determine a real withdrawal rate such that half of the real wealth at the start of retirement remained after 20 years, with high probability. Rather than using fixed portfolio weights as in Bengen (1994), the portfolio was invested in stocks and bonds according to the QS optimal strategy described above. Dang et al. concluded that the 4% rule still held up well under the revised assumptions.

An obvious alternative way to generate retirement cash flows is to buy a lifetime annuity. However, in practice most retirees are not willing to do this, for a variety of reasons (MacDonald et al., 2013). In the current environment of low real interest rates, annuities provide rather low income, so the reluctance of retirees to use them is particularly unsurprising.

Consequently, we pose the pension accumulation problem as follows. The desired expected accumulated real wealth at retirement W_d is

$$E[W_T] = W_d = \frac{R e^{\mu_I T} I_0}{w_r}, \quad (4.1)$$

where R is the replacement ratio, I_0 is the initial salary, μ_I is the real salary escalation rate, T is the end of the accumulation period, and w_r is the safe withdrawal rate. Recall that above we denoted the set of action times prior to T as \mathcal{T}_1 . We assume that cash is contributed into the portfolio at time $t_i \in \mathcal{T}_1$. The initial cash contribution at t_0 is $F_c I_0$, where F_c is the real contribution fraction. This contribution fraction represents the total amount contributed by both the employee and the employer to all retirement savings accounts, but excludes any government sponsored universal schemes (e.g. CPP in Canada, Social Security in the US). We also assume that these accounts are tax-advantaged, i.e. no tax is paid during the accumulation phase. At subsequent action times $t_i \in \mathcal{T}_1$ (i.e. after the initial contribution), the amount contributed is assumed to be $F_c I_0 e^{\mu_I t_i}$.

Table 4.1 summarizes the data for this base case scenario. As indicated in the table, we assume a replacement fraction of 50%, an initial salary of \$50,000, and a real escalation rate of 1.27%.¹¹ Combined contributions by the employee and employer to the retirement savings portfolio are 20%

¹¹As noted by Bloom et al. (2014), this rate has been used by the US Congressional Budget Office in its long-term projections.

Initial salary I_0	\$50,000
Salary escalation rate μ_I	0.0127 (Bloom et al., 2014)
Contribution fraction F_c	0.20
Accumulation period T	30 years
Safe withdrawal rate w_r	0.04
Equity index	Value-weighted
Bond index	3-month T-bill
Investment strategies	Constant proportion, glide path, QS optimal
Rebalancing interval	1 year
Maximum leverage ratio L_{\max}	1.0
If insolvent	Trading stops

TABLE 4.1: *Data for base case scenario. Cash is injected into the portfolio at times $t = 0, 1, \dots, 29$. Market parameters for the equity and bond indexes are provided in Table 3.1.*

of real salary each year. For simplicity, these contributions are assumed to be made at the start of each year during a 30 year accumulation period. The safe withdrawal rate is assumed to be 4%, in line with Bengen (1994) and Dang et al. (2017). Applying these parameters to equation (4.1), we find that the expected (real) desired terminal wealth W_d is

$$E[W_T] = W_d = \frac{.50}{.04} \times \$50,000 e^{.0127 \times 30} \simeq \$915,000. \quad (4.2)$$

Note that the specification given in equation (4.1) implies that decreasing the withdrawal rate w_r has the same effect as increasing the replacement fraction R or the salary escalation rate μ_I .

As indicated in Table 4.1, the retirement savings portfolio is invested in the value-weighted index and the 3-month T-bill index. Relevant parameters for these indexes are given in Table 3.1. We consider three alternative investment strategies:

- Constant proportion, i.e. $p = \text{const.}$
- Linear glide path, as in equation (2.5).
- Time consistent QS optimal strategy, as described in Section 2.2. Recall that this strategy is also multi-period pre-commitment MV optimal.

Each of these strategies is rebalanced annually. We start with the constant proportion strategy, and determine the equity weight such that $E[W_T] = \$915,000$, assuming the market parameters given in Table 3.1 for the value-weighted equity index and the 3-month T-bill index. This turns out to be 0.5788.¹² We then turn to the linear glide path strategy. In this case we specify $p_{\max} = 1.0$, and then determine the value of $p_{\min} = 0.3066$ is needed to have $E[W_T] = \$915,000$. We proceed similarly for the QS optimal strategy given by the solution to problem (2.8). Imposing the leverage constraint $L_{\max} = 1.0$, we find by Newton iteration that the value of W^* which results in $E[W_T] = \$915,000$ is \$1,106,200. We compute and store the optimal control associated with this value of W^* in order to test this strategy. Note that:

¹²This is a bit more aggressive in terms of taking on equity market risk than the strategy considered by Bengen (1994) which involved equal weights between the equity and bond markets. Keep in mind that here we are investing in a 3-month T-bill index, whereas Bengen used intermediate maturity Treasury bonds which offer somewhat higher average returns.

- When we change input parameters (e.g. invest in different assets, allow $L_{\max} > 1$, etc.), we may need to recompute the expected wealth target $E[W_T] = W_d$, the equity weight for the constant proportion strategy, the glide path parameters (p_{\max}, p_{\min}) , and the quadratic wealth target W^* (along with the associated optimal control) in order to meet this target.
- The quadratic wealth target W^* exceeds the target expected real terminal wealth $W_d = E[W_T]$. This is because the QS optimal strategy will de-risk if W^* is attainable by investing only in the risk-free asset, so there is not much chance of exceeding this quadratic target by a significant amount. This implies that the average terminal wealth, factoring in paths where the accumulated savings does not ever reach W^* , must be lower than W^* .

We compare these three strategies using two different types of simulation tests. As an initial test, we assume that the stochastic environment described in Table 3.1 holds exactly. In other words, the level of the equity market index follows a double exponential jump diffusion with the parameters given in Table 3.1 and the bond market index is non-stochastic, with a constant risk-free interest rate as indicated in that table. We refer to this as a *synthetic market*. In such a market, we draw 160,000 Monte Carlo simulated paths and compute performance statistics. Note that these comparisons are based on a simulated environment that corresponds exactly to the environment used to formulate the strategies. As a second and more stringent test, we draw simulated paths by bootstrap resampling of the historical return data and compute the same performance statistics. We refer to this type of backtest as a *historical market*. This is a stricter test since it does not assume that the equity market follows a jump diffusion process or that the risk-free interest rate is constant over time, although those assumptions are still used to generate the strategies that are followed. A single resampled path is constructed by pasting together enough blocks of monthly historical return data to cover the investment horizon of 30 years. The sampling is done in blocks to account for possible serial dependence. The blocks are selected simultaneously from both the historical stock and bond market indexes, to incorporate possible correlations. The blocks are chosen randomly, with replacement. To avoid end issues, the historical data is wrapped around.¹³ To reduce the impact of a fixed blocksize and mitigate edge effects at each block end, we use the stationary block bootstrap (Politis and White, 2004; Patton et al., 2009) where the blocksize is sampled randomly from a geometric distribution with an expected value of \hat{b} . In principle, the optimal expected blocksize can be estimated using an algorithm provided by Patton et al. (2009). As discussed in Forsyth and Vetzal (2019), this approach is not easily applied in our context. This is because the estimated optimal blocksizes for the different market indexes we consider (i.e. the value-weighted and equal-weighted US equity market indexes and the 3-month T-bill and 10-year T-bond indexes) vary considerably, ranging from about two months for the value-weighted stock index to more than four years for the T-bill index. Recall that we sample simultaneously from both a stock index and a bond index, so we must use the same blocksize for both indexes, and our strategies involve weighted combinations of two of these indexes that can change deterministically (glide path) and also randomly (QS optimal) along a simulation path. As a result, we report results for a range of expected blocksizes \hat{b} , acknowledging that the choice of \hat{b} for our application is open to debate.

Table 4.2 gives the results for the base case input data from Table 4.1. Consider first the synthetic market. By construction, all three strategies have the same expected value of real terminal wealth at the retirement date of 915 (units in the table are in thousands of dollars). The constant proportion and glide path strategies are effectively indistinguishable, having the same standard

¹³In other words, if the size of a block extends past the end of the sample in 2015:12, the return data resumes at the start of the sample in 1926:1 for the duration of the block.

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 700)$	$Pr(W_T < 800)$	Expected Surplus Cash
Synthetic Market					
Constant proportion	915	519	.39	.51	NA
Glide path	915	519	.39	.51	NA
QS optimal	915	244	.19	.24	21
Historical Market (Expected blocksize $\hat{b} = 1$ year)					
Constant proportion	876	402	.38	.51	NA
Glide path	872	398	.39	.52	NA
QS optimal	904	232	.18	.24	26
Historical Market (Expected blocksize $\hat{b} = 2$ years)					
Constant proportion	869	376	.38	.51	NA
Glide path	866	372	.39	.52	NA
QS optimal	911	221	.17	.23	31
Historical Market (Expected blocksize $\hat{b} = 5$ years)					
Constant proportion	862	349	.37	.50	NA
Glide path	861	347	.38	.51	NA
QS optimal	924	213	.16	.21	38

TABLE 4.2: *Base case scenario results. Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12.*

deviation of terminal wealth and the same probability of ending up with wealth below 700 or below 800. By contrast, the QS optimal strategy has much lower standard deviation or shortfall probability for those two levels of terminal wealth, by a factor of around two in each case. In addition, this strategy offers a small amount of expected surplus cash (this surplus is not applicable for the other two strategies). Turning to the historical market results for expected blocksizes of 1, 2, and 5 years, we reach generally similar conclusions. Of course, the expected values no longer are exactly equal to the target of 915, but the difference from this target is much lower for the QS optimal strategy. The constant proportion and glide path strategies are again quite comparable to each other, with the glide path having slightly lower expected value and standard deviation, but a bit higher shortfall probability (for the two values of W_T considered). In all cases, the QS optimal strategy offers the best performance, with higher mean, lower standard deviation, and lower shortfall probability, as well as a modest amount of expected surplus cash.

Figure 4.1(a) shows the cumulative distributions for all three strategies in the historical market (expected blocksize of two years). The distributions for the constant proportion and glide path strategies are practically identical. Over the bulk of the distribution, the QS optimal strategy exhibits clearly better performance than the others. However, the QS optimal strategy performs a little worse than the alternatives in the extreme left tail. This happens because there can be paths where the equity market trends downward for a very large portion of the investment horizon. In such cases, all strategies do poorly, but the QS optimal strategy will remain fully invested in the equity market in an attempt to recover and meet the quadratic wealth target. The QS optimal strategy also underperforms in the extreme right tail of the distribution. This is because there are paths where the equity market trends strongly upward over most of the investment period. Once the quadratic wealth target is reached, however, the QS optimal strategy de-risks, shifting all investment into the

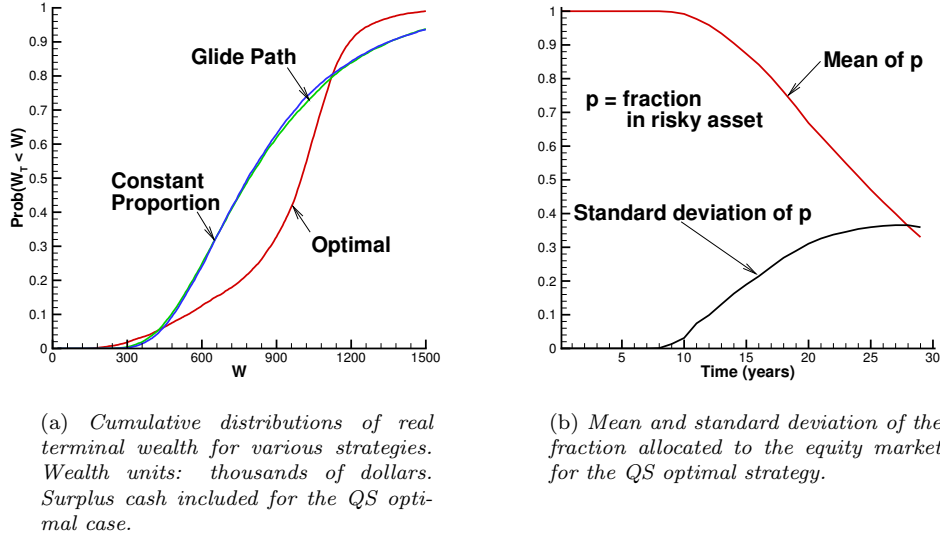


FIGURE 4.1: Base case scenario results in the historical market. Input data provided in Tables 3.1 and 4.1. Results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12 with expected blocksize $\hat{b} = 2$ years.

low return bond market. It does not capitalize on the continued strong equity market performance. The other two strategies, by comparison, retain a large equity market exposure, leading to higher terminal wealth. However, we reiterate that over most of the distribution, the QS optimal strategy provides better results.

Figure 4.1(b) depicts properties of the optimal control for the QS optimal strategy. This strategy invests entirely in the equity market for the first several years. The percentage invested in the risky asset subsequently trends downward on average over time. However, there is considerable variation: the standard deviation of the optimal control rises strongly over time, indicating that the allocation to equities is quite sensitive to realized investment returns.¹⁴

The historical market results given in Table 4.2 are somewhat encouraging for the QS optimal strategy. To formulate this strategy, recall that we assumed a double exponential jump diffusion model with known parameters for the equity market index and a constant risk-free interest rate. Our bootstrap resampled historical market tests make no such assumptions, yet deliver results that are fairly close to those observed in the synthetic market tests which maintain these assumptions. This indicates that the QS optimal strategy is quite robust to departures from these assumptions.

As an additional test of the robustness of the QS optimal strategy, we explore the effect of computing and storing the optimal control based on constant parameters from Table 3.1 but then allowing the synthetic equity market parameters to vary randomly in simulation tests. To be specific, we carry out Monte Carlo simulations where at each action time $t_i \in \mathcal{T}_1$ and along each stochastic path we select (μ, σ) from a uniform distribution having mean equal to the corresponding values from Table 3.1. This set of (μ, σ) is then used for the interval (t_i, t_{i+1}) .

Table 4.3 shows the results. The first row reproduces the values reported in Table 4.2 for this strategy in the synthetic market, while the remaining two rows provide results when μ and σ are varied randomly within the given range. Table 4.3 indicates that only an estimate of the mean

¹⁴Of course, the equity allocation for the constant proportion and glide path cases is fixed in advance, being at most time-dependent and not varying at all in response to realized returns.

Market Parameters	Expected Value	Standard Deviation	$Pr(W_T < 700)$	$Pr(W_T < 800)$	Expected Surplus Cash
Synthetic market	915	244	.19	.24	21
$\mu \in [.04889, .1289]$	916	245	.19	.24	21
$\sigma \in [0.1077, 0.1877]$	915	245	.19	.24	21

TABLE 4.3: *Base case scenario results for the QS optimal strategy with random variation of market parameters μ and σ , Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1, except as noted. Monte Carlo simulations with (μ, σ) drawn from a uniform distribution with the indicated limits along each of 160,000 paths.*

of the distribution of the market parameters μ and σ is needed to compute an effective control strategy. This is consistent with the results in Ma and Forsyth (2016), where it is shown that including stochastic volatility effects results in negligible improvements in results for long term (i.e. greater than ten years) investors.

Overall, then, for the base case data in Table 4.1, the QS optimal strategy appears to be fairly robust to parameter and model mis-specification. Moreover, this strategy clearly outperforms the constant proportion and glide path alternatives over most of the terminal wealth distribution. It provides about the same mean terminal wealth, but has considerably lower standard deviation and shortfall probability (for the two wealth levels considered in Table 4.2). However, even the QS optimal strategy delivers somewhat disappointing results in absolute terms. Recall that we are trying to achieve average real terminal wealth of \$915,000. In the idealized synthetic market which conforms exactly to the modelling assumptions used to generate the controls for the QS optimal strategy, Table 4.2 shows that there is almost a 20% chance of ending up with real terminal wealth below \$700,000, and about a 25% chance of ending up with less than \$800,000. These shortfall probabilities do not change much under the conditions of the historical market backtests. Of course, these rather pessimistic results so far have only considered the base case data. In the following section, we investigate whether more promising results can be achieved under different assumptions.

5 Alternative Assumptions

The base case results in Section 4 relied on the input parameters given in Tables 4.1 and 3.1, for the value-weighted equity and 3-month T-bill indexes. We now explore the effects of altering our assumptions about factors such as the contribution fraction, the salary escalation rate, the maximum amount of leverage permitted, the underlying indexes to be used, and the salary replacement ratio. We consider each of these in turn. For the most part, we only use the QS optimal strategy since it has been shown above to be generally superior to the constant proportion and glide path strategies.

5.1 Effect of Contribution Fraction

Our base case described by Table 4.1 assumed a total combined contribution by the employee and employer of $F_c = 20\%$ of salary. Table 5.1 reports the effects of dropping this to 15% or increasing it to 25% for the QS optimal strategy. As is to be expected, the table shows that risk (measured either in terms of standard deviation or the reported shortfall probabilities) decreases significantly as F_c rises. However, even in the case where 25% of the employee's salary is contributed to the retirement savings plan, there is still almost a 15% chance that real terminal wealth is less than

F_c	Expected Value	Standard Deviation	$Pr(W_T < 700)$	$Pr(W_T < 800)$	Expected Surplus Cash
Synthetic Market					
0.15	915	440	.36	.42	18
0.20	915	244	.19	.24	21
0.25	915	150	.09	.13	21
Historical Market (Expected blocksize $\hat{b} = 2$ years)					
0.15	916	380	.34	.41	12
0.20	911	221	.17	.23	31
0.25	909	126	.07	.14	67

TABLE 5.1: *Effect of varying contribution fraction F_c for the QS optimal strategy. Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12.*

\$800,000, considerably lower than the target of \$915,000. A broader comparison of these cases is provided in Figure 5.1, which depicts the cumulative distributions of real terminal wealth. The cases where $F_c = 20\%$ and $F_c = 25\%$ appear quite comparable for values of $W_T \geq \$915,000$ over the plotted range, but the higher contribution fraction case appears to be much safer over a wide range below the target. The high savings rate leads to a notably increased amount of expected surplus cash. Along paths with strong equity market returns, the target can be reached relatively early, but the model assumes that savings continue each year (in this case at a high rate), so surplus cash can build up. The case where $F_c = 15\%$ exhibits poor performance on the downside, but is somewhat better on the upside. With a low amount of money saved, more risk must be taken on in order to reach the target. Doing so works out very well if realized returns are strong, and quite poorly if they are not.

5.2 Effect of Salary Escalation Rate

We now examine the effect of changing the salary escalation rate μ_I from the base case value of 1.27% given in Table 4.1. Note that this will result in a different expected real wealth target W_d , based on equation (4.1). Table 5.2 summarizes the results for both the synthetic market and the historical market, using the QS optimal control. For comparability, the table expresses the results in terms of W_T/W_d , i.e. real terminal wealth as a fraction of the expected wealth target. Obviously, a higher escalation rate leads to a higher final salary. Given a fixed salary replacement ratio, this translates into a higher expected value of terminal wealth W_d . For example, with $\mu_I = 1.75\%$ we have $W_d = \$1,056,000$ instead of \$915,000 as in the base case. The results in Table 5.2 are quite similar in the synthetic and historical markets, on a case by case basis. For each set of tests, the standard deviation and shortfall probabilities show an increase with μ_I . With a higher salary and a fixed contribution fraction, there will obviously be a higher amount of saving. Despite this, the associated higher real terminal wealth target results in a higher level of risk. This is borne out in Figure 5.2, which plots the cumulative distribution functions of normalized real terminal wealth W_T/W_d for the various values of μ_I . The highest salary escalation rate has the worst performance for low W_T/W_d , and the best performance for high W_T/W_d . Taking on more risk to reach the higher wealth target works out well if investment returns are favourable, and poorly if they are not. Conversely, the lowest value of μ_I results in the best performance if investment returns are weak,




FIGURE 5.1: *Cumulative distributions of real terminal wealth for various contribution fractions F_c for the QS optimal strategy. Wealth units: thousands of dollars. Input data provided ‘in Tables 3.1 and 4.1, except as noted. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12 with expected blocksize $\hat{b} = 2$ years. Surplus cash included.*

and the worst performance if they are not.

5.3 Effect of Leverage

All of the results presented thus far have specified $L_{\max} = 1$, as indicated in Table 4.1. We now consider increasing this to $L_{\max} = 1.5$, thereby allowing the use of 50% leverage. Table 5.3 documents the results for the QS optimal strategy. It is interesting to observe that the risk measures shown (standard deviation of W_T and the two shortfall probabilities) each indicate somewhat lower risk when the use of leverage is permitted. In one sense, this is obvious: relaxing a constraint cannot lead to worse performance. On the other hand, leverage is often not allowed since it is perceived to be “risky”. The issue here is what is meant by risk. We have defined it on the basis of the value of terminal wealth. Leverage constraints, however, are typically motivated by concerns that portfolio values during the accumulation period (not at the end of the period) will fluctuate excessively.

Figure 5.3(a) shows the cumulative distributions of real terminal wealth when leverage is allowed and when it is not. The cumulative distributions are fairly similar for high levels of terminal wealth. Over a wide range of wealth levels below the target of \$915,000, the strategy which permits leverage performs better. However, in the extreme left tail of the distribution, it turns out to be worse to allow leverage. These very low wealth levels occur as a result of very poor equity market returns over most of the investment horizon. The QS optimal strategy continues to try to reach the quadratic wealth target, so it invests completely in the equity market to the extent possible. With continued poor returns, leverage in this case amounts to doubling down on a losing bet.

μ_I	Wealth Target W_d	Expected Value (W_T/W_d)	Standard Deviation (W_T/W_d)	$Pr(W_T/W_d < 0.8)$	$Pr(W_T/W_d < 0.9)$	Expected Surplus Cash ($/W_d$)
Synthetic Market						
0.0175	1056	1.0	.31	.25	.30	.02
0.0127	915	1.0	.27	.20	.25	.02
0.0075	783	1.0	.22	.15	.21	.02
Historical Market (Expected blocksize $\hat{b} = 2$ years)						
0.0175	1056	.99	.29	.23	.29	.03
0.0127	915	.97	.24	.18	.25	.03
0.0075	783	.99	.19	.14	.21	.04

TABLE 5.2: *Effect of varying salary escalation rate μ_I for the QS optimal strategy. Units for W_d : thousands of dollars. Remaining wealth values are normalized by W_d for each case. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12.*

L_{\max}	Expected Value	Standard Deviation	$Pr(W_T < 700)$	$Pr(W_T < 800)$	Expected Surplus Cash
Synthetic Market					
1.0	915	244	.19	.24	21
1.5	915	205	.12	.17	24
Historical Market (Expected blocksize $\hat{b} = 2$ years)					
1.0	911	221	.17	.23	31
1.5	904	186	.11	.18	46

TABLE 5.3: *Effect of varying maximum leverage ratio L_{\max} for the QS optimal strategy. Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12.*

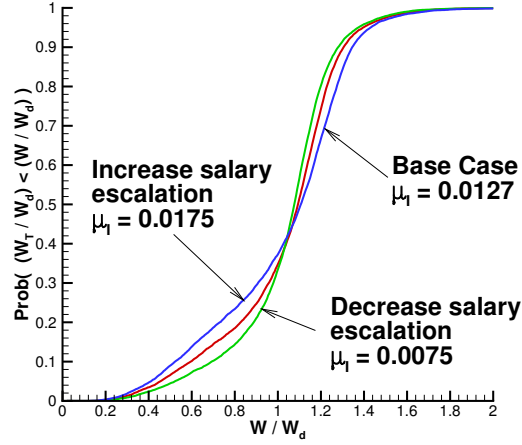


FIGURE 5.2: Cumulative distributions of normalized real terminal wealth for various salary escalation rates μ_I for the QS optimal strategy. Input data provided in Tables 3.1 and 4.1, except as noted. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12 with expected blocksize $\hat{b} = 2$ years. Surplus cash included.

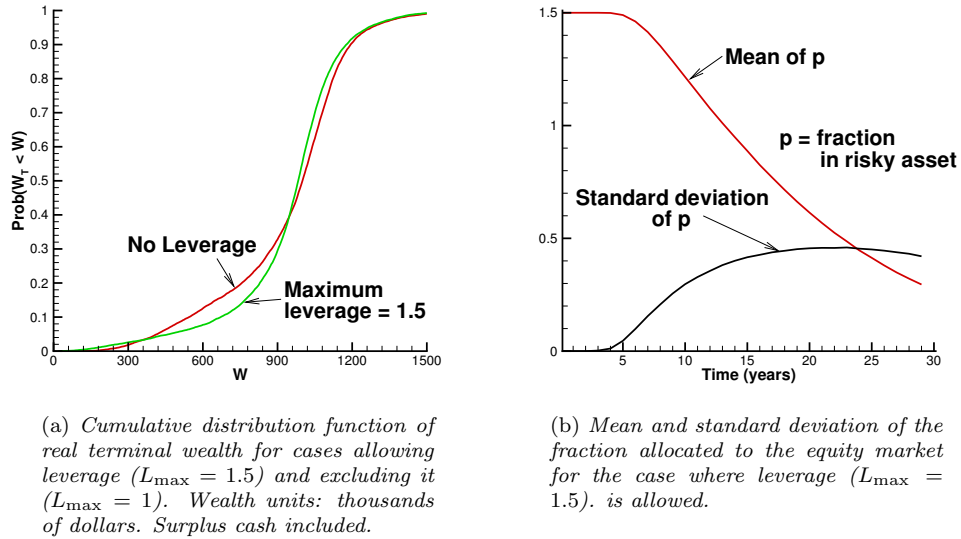


FIGURE 5.3: Results for cases allowing and excluding leverage. Input data provided in Tables 3.1 and 4.1, except as noted. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12 with expected blocksize $\hat{b} = 2$ years.

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 700)$	$Pr(W_T < 800)$	Expected Surplus Cash
Synthetic Market					
Constant proportion	915	437	.34	.47	NA
Glide path	915	438	.34	.48	NA
QS optimal	915	222	.16	.21	19
Historical Market (Expected blocksize $\hat{b} = 2$ years)					
Constant proportion	900	374	.34	.47	NA
Glide path	897	376	.34	.48	NA
QS optimal	881	201	.17	.26	67

TABLE 5.4: *Results for the optimal QS strategy when the 10-year T-bond index is used instead of the 3-month T-bill index. Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12.*

5.4 Long Term Bond Index

All of the results presented thus far were based on the value-weighted equity index and the 3-month T-bill index. In this section we explore the effects of substituting the 10-year T-bond index for the T-bill index. Recall from Table 3.1 that the long term bond index had average annual real returns of 2.16%, about 1.3% higher than the T-bill index. Of course, these higher average returns are accompanied by higher volatility (Figure 3.1), but that is ignored when we formulate the QS optimal strategy since we assume constant interest rates. Table 5.4 shows results for the constant proportion, linear glide path, and QS optimal strategies when the 10-year T-bond index is used in lieu of the 3-month T-bill index. The general pattern in Table 5.4 is similar to that seen earlier in Table 4.2. The QS optimal approach achieves roughly the same expected real terminal wealth, but takes on considerably lower risk compared to the other two strategies, as measured by standard deviation of W_T or either of the two shortfall probabilities considered. Comparing the results from Tables 4.2 and 5.4, we observe that the standard deviation and shortfall probabilities for each strategy are a little lower in the synthetic market when the 10-year T-bond index is used. This also holds in the historical market for the constant proportion and glide path strategies, but is not the case for the QS optimal strategy in terms of the shortfall probabilities. For example, the probability of real terminal wealth being lower than \$800,000 is 23% if the 3-month T-bill index is used, but 26% when the 10-year T-bond index is used. The overall conclusion, however, is that replacing the 3-month T-bill index by the 10-year T-bond index does not make much difference. This picture is reinforced by comparing the cumulative distributions of real terminal wealth shown in Figure 5.4 based on the 10-year T-bond index with those shown previously in Figure 4.1(a) for the 3-month T-bill index.

5.5 Equal-Weighted Equity Index

As shown above in Figure 3.2, an equal-weighted equity index has historically outperformed the value-weighted equity index that has been used for all results to this point. We now consider the impact of replacing the value-weighted index by its equal-weighted counterpart. Note that the bond index used here is the 3-month T-bill index. Table 5.5 provides the results. Whether in the idealized synthetic market or the backtest historical market, the QS optimal strategy clearly

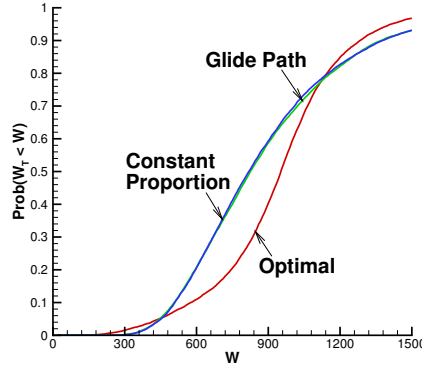


FIGURE 5.4: Cumulative distribution of real terminal wealth for QS optimal, glide path, and constant proportion strategies when the 10-year T-bond index is used. Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1, except as noted. Surplus cash included for the QS optimal strategy. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12. Expected blocksize $\hat{b} = 2$ years.

outperforms the constant proportion and glide path alternatives by the criteria given in the table. In the synthetic market, the QS optimal strategy achieves the same $E[W_T]$ with dramatically lower standard deviation of W_T and shortfall probability, along with the possibility of a modest amount of surplus cash. The same general conclusions apply in the historical market, although it is worth noting that the average real terminal wealth for the other strategies is somewhat lower for all expected blocksizes considered. Comparing the results in Table 5.5 for the equal-weighted equity index with those reported above in Table 4.2 for the value-weighted index, it can be seen that the shortfall probabilities are now considerably lower for the QS optimal strategy, but almost unchanged for the other strategies. The standard deviation of W_T , however, is substantially lower for all of the strategies.

Figure 5.5 plots the cumulative distributions of real terminal wealth in both the synthetic and the historical markets. In both cases, the distributions for the glide path and constant proportion strategies are virtually indistinguishable. Figure 5.5(a) indicates that the QS optimal strategy outperforms over a wide range of terminal wealth values in the synthetic market, although it does perform worse in the tails of the distribution. As remarked earlier, this is due to two features of the strategy: (i) it automatically de-risks once the quadratic wealth target is achievable by investing only in the bond index (so it does not take advantage of continued strong equity market performance afterwards on paths where that happens); and (ii) it continually tries to reach the quadratic wealth target by using maximum equity market exposure (and this gamble for resurrection fails on paths where the equity market has persistently poor performance). The same comments apply to the historical market shown in Figure 5.5(b), but it is worth noting that the underperformance of the QS optimal strategy in the tails is considerably reduced here compared to the synthetic market.

5.6 Effect of Replacement Ratio

All of the results provided to here assume a replacement ratio $R = 50\%$ of final real salary, in accordance with Table 4.1. We now explore the effects of lowering this to 40% and increasing it to 60%. For each case, we determine the desired expected real wealth target by using equation (4.1). Table 5.6 shows the results. Decreasing R to 40% reduces the expected wealth target W_d to

Strategy	Expected Value	Standard Deviation	$Pr(W_T < 700)$	$Pr(W_T < 800)$	Expected Surplus Cash
Synthetic Market					
Constant proportion	915	546	.38	.51	NA
Glide path	915	553	.39	.52	NA
QS optimal	915	185	.11	.16	44
Historical Market (Expected blocksize $\hat{b} = 1$ year)					
Constant proportion	837	327	.39	.54	NA
Glide path	831	319	.39	.55	NA
QS optimal	904	162	.10	.16	49
Historical Market (Expected blocksize $\hat{b} = 2$ years)					
Constant proportion	827	293	.38	.54	NA
Glide path	820	283	.38	.55	NA
QS optimal	915	139	.07	.13	51
Historical Market (Expected blocksize $\hat{b} = 5$ years)					
Constant proportion	815	248	.36	.53	NA
Glide path	808	242	.37	.55	NA
QS optimal	932	115	.04	.10	54

TABLE 5.5: *Results when the equal-weighted equity index is used instead of the value-weighted equity index. Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12.*

\$732,000 from \$915,000, while raising R to 60% increases W_d to \$1,098,000. The remaining wealth values in the table are normalized by W_d . Whether we consider the synthetic or the historical market, it is clear that increasing R requires taking on more risk, as measured by either the standard deviation or the shortfall probabilities. This is borne out in the cumulative distribution plots of normalized terminal real wealth provided in Figure 5.6. The synthetic market results in Figure 5.6(a) clearly indicate that the QS optimal strategy performs better for the lowest value of R . Recall that the strategy attempts to come as close as possible to W_d (normalized to 1 for this plot), and the cumulative distribution when $R = 40\%$ shows relatively low probability of normalized real wealth being much above or below 1. For the highest replacement ratio ($R = 60\%$), there is a substantial chance of being either significantly below or above 1. This is because the strategy must take on more risk in order to attain the higher expected wealth target. Of course, the base case with $R = 50\%$ lies in between these other two cases. The results for the historical market shown in Figure 5.6(b) are generally similar, though the differences across the range of values of R are somewhat less pronounced.

5.7 Summary Regarding Alternative Assumptions

Sections 5.1-5.6 above provided detailed results concerning the effects of

- varying the accumulation fraction F_c , i.e. the investor saves $F_c I_0 e^{\mu_I t_i}$ at each rebalancing date;
- varying the real salary escalation rate μ_I ;

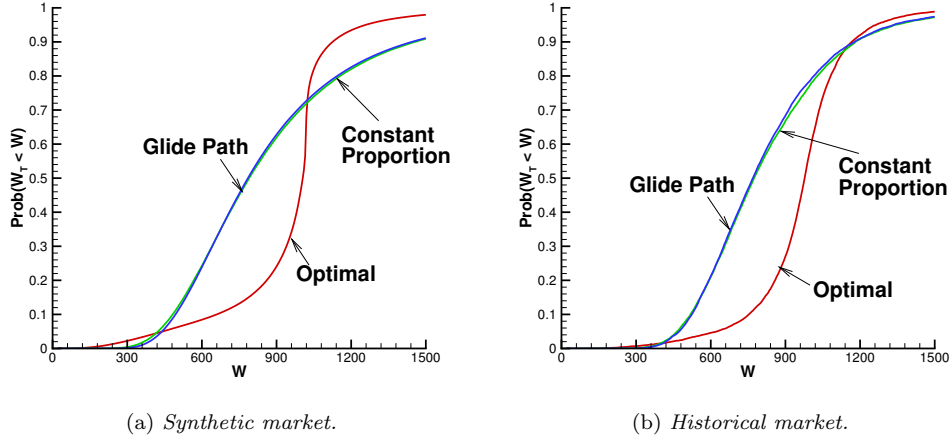


FIGURE 5.5: Cumulative distribution of real terminal wealth when the equal-weighted equity index is used. Wealth units: thousands of dollars. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12 with expected blocksize $\hat{b} = 2$ years. Surplus cash flow included for the QS optimal strategy.

- use of leverage;
- alternative bond index: use of a 10 year T-bond index instead of a 3 month T-bill index;
- alternative stock index: use of an equal-weighted equity index instead of a value-weighted index; and
- varying the replacement ratio R .

Table 5.7 summarizes the results of various strategies in terms of probability of shortfall with respect to the desired wealth goal W_d . These results were all obtained using bootstrap resampling (i.e. for the historical market) with an expected blocksize of two years. As a filter to determine an acceptable combination of DC plan parameters and investment strategies, suppose we specify that there should be at least a 90% probability of achieving at least 80% of the desired expected wealth goal W_d . Based on attempting to achieve the final target expected real wealth for the base case (see Table 4.1) and applying this filter, we can see that the shortfall probabilities using standard strategies (constant proportion or glide path) are unacceptably high. Using the QS optimal strategy leads to a substantial reduction in these shortfall probabilities, but still not to the desirable range of less than 10%.

From equation (4.1), it is clear that the case with $R = .6$ and $w_r = 4\%$ leads to the same expected wealth target as specifying $R = .5$ and $w_r = 3.3\%$. Table 5.7 therefore indicates that if we assume that the safe (real) withdrawal rate is 3.3% and the replacement ratio is 50%, the probability of shortfall is quite high even if the QS optimal strategy is followed.

Assuming we use the QS optimal allocation strategy, we are then forced to take other actions to attempt to reduce the shortfall probability. Increasing the contribution rate to 25% of annual salary meets our criterion, but this might be difficult to implement in terms of agreement from employees and employers. Decreasing the replacement ratio (40% of final salary) also achieves the shortfall objective. We note that many institutions effectively do this by targeting a final career average salary replacement ratio (instead of a final salary replacement ratio).

R	Wealth Target W_d	Expected Value (W_T/W_d)	Standard Deviation (W_T/W_d)	$Pr(W_T/W_d < 0.8)$	$Pr(W_T/W_d < 0.9)$	Expected Surplus Cash ($/W_d$)
Synthetic Market						
0.4	732	1.0	.16	.10	.15	.02
0.5	915	1.0	.27	.20	.25	.02
0.6	1098	1.0	.37	.31	.36	.02
Historical Market (Expected blocksize $\hat{b} = 2$ years)						
0.4	732	.99	.15	.09	.17	0.0
0.5	915	.97	.24	.18	.25	.03
0.6	1098	1.0	.35	.30	.35	.02

TABLE 5.6: *QS optimal results with varying salary replacement ratios R . Units for W_d : thousands of dollars. Remaining wealth values are normalized by W_d for each case. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12.*

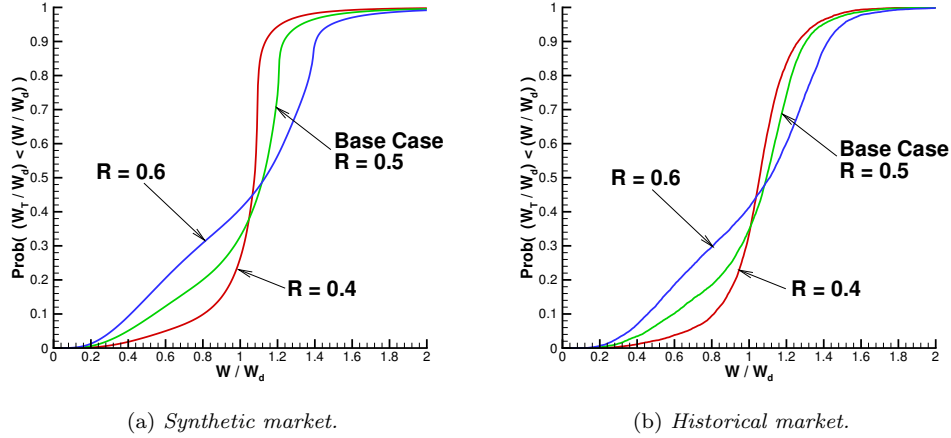


FIGURE 5.6: *Cumulative distributions of real terminal wealth with different salary replacement ratios R for the QS optimal strategy. Input data provided in Tables 3.1 and 4.1, except as noted. Synthetic market results computed using Monte Carlo simulations with 160,000 sample paths. Historical market results based on 10,000 bootstrap resampled paths using historical data from 1926:1 to 2015:12 with expected blocksize $\hat{b} = 2$ years. Surplus cash included.*

Case	$Pr(W_T/W_d < 0.8)$	$Pr(W_T/W_d < 0.9)$
Base Case, Table 4.1		
Constant proportion	.43	.54
Glide path	.43	.54
QS optimal	.18	.25
QS Optimal		
Contribution Fraction		
$F_c = .25$.09	.17
$F_c = .15$.37	.43
Salary Escalation Rate		
$\mu_I = .0175$.23	.29
$\mu_I = .0075$.14	.21
Leverage		
$L_{\max} = 1.5$.13	.19
Alternative Bond Index		
10-year T-bond	.19	.28
Alternative Stock Index		
Equal-weighted	.09	.16
Replacement Ratio		
$R = .4$.09	.17
$R = .6$.30	.35

TABLE 5.7: *Comparison of shortfall probabilities. Results are normalized by W_d for each case. Input data provided in Tables 3.1 and 4.1, except as noted. Historical market results based on 10,000 bootstrap resampled paths using data from 1926:1 to 2015:12 with expected blocksize $\hat{b} = 2$ years.*

Finally, the use of the alternative equal-weighted equity index also achieves the shortfall probability target. As noted earlier, this type of index has historically outperformed its value-weighted counterpart owing to higher exposure to value, size, and market factors (Plyakha et al., 2014). In effect, the equal-weighted portfolio is a *smart beta* portfolio, with a long track record. However, equal-weighted portfolios have higher costs, which have not been factored in to our analysis. This suggests that there may be a market opportunity for a low cost synthetic ETF which tracks the equal-weighted index.

6 Conclusions

Typical default strategies available to DC plan members include target date funds (i.e. glide path) and balanced funds (constant proportion). Both of these strategies have high probabilities of shortfall and large standard deviations of real terminal wealth.

We define an acceptable probability of success as a 90% probability of achieving 80% of the real desired terminal wealth goal. QS optimal strategies reduce the probability of shortfall significantly, compared to the ubiquitous glide path and balanced portfolio strategies, but not to within an acceptable range.

Assuming an optimal QS strategy, acceptable probabilities of shortfall can be obtained by

- reducing the final salary target replacement ratio (40% or less);
- increasing the total (employee and employer) contribution rate to 25% per year;
- using alternative stock investment indices, such as an equal-weighted index. The backtests of an equal-weighted index perform well, but it is not clear that this will persist in the future. In addition, we have not factored in the additional costs of this type of index.

Our main conclusion is that current practices in DC plans with typical target benefits and default investment strategies have unacceptably high probabilities of shortfall. This has major ramifications for the many organizations which are shifting to DC plans.

Our results are based on the past 90 years of US market data. It is possible to argue that future market returns will be lower than observed in the past. The implication of such an assumption is that the probability of success for a DC investor would be even lower than we have reported here. In this case, the situation for DC plan investors would be even more dire than we suggest.

Acknowledgements

P. A. Forsyth's work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) grant RGPIN-2017-03760.

Conflicts of interest

The authors have no conflicts of interest to report.

References

- Arnott, R. D., K. F. Sherrerd, and L. Wu (2013). The glidepath illusion and potential solutions. *The Journal of Retirement* 1(2), 13–28.
- Barber, B. M. and T. Odean (2013). The behavior of individual investors. In G. Constantinides, M. Harris, and R. Stulz (Eds.), *Handbook of Economics and Finance*, Chapter 22, pp. 1533–1569. Elsevier.
- Basak, S. and G. Chabakauri (2010). Dynamic mean-variance asset allocation. *Review of Financial Studies* 23, 2970–3016.
- Basu, A. K., A. Byrne, and M. E. Drew (2011). Dynamic lifecycle strategies for target date retirement funds. *Journal of Portfolio Management* 37(2), 83–96.
- Bengen, W. (1994). Determining withdrawal rates using historical data. *Journal of Financial Planning* 7, 171–180.
- Bensoussan, A., K. C. Wong, and S. C. P. Yam (2019). A paradox in time-consistency in the mean-variance problem? *Finance and Stochastics* 23, 173–207.
- Bjork, T. and A. Murgoci (2010). A general theory of Markovian time inconsistent stochastic control problems. SSRN 1694759.
- Bjork, T. and A. Murgoci (2014). A theory of Markovian time inconsistent stochastic control in discrete time. *Finance and Stochastics* 18, 545–592.
- Bjork, T., A. Murgoci, and X. Y. Zhou (2014). Mean-variance portfolio optimization with state-dependent risk aversion. *Mathematical Finance* 24, 1–24.
- Bloom, D. E., D. Canning, and M. Moore (2014). Optimal retirement with increasing longevity. *Scandinavian Journal of Economics* 116, 838–858.
- Cocco, J. F., F. J. Goems, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies* 18, 491–533.
- Cogneau, P. and V. Zakalmouline (2013). Block bootstrap methods and the choice of stocks for the long run. *Quantitative Finance* 13, 1443–1457.
- Cont, R. and C. Mancini (2011). Nonparametric tests for pathwise properties of semimartingales. *Bernoulli* 17, 781–813.
- Cui, X., J. Gao, X. Li, and D. Li (2014). Optimal multi-period mean variance policy under no-shorting constraint. *European Journal of Operational Research* 234, 459–468.
- Cui, X., D. Li, S. Wang, and S. Zhu (2012). Better than dynamic mean-variance: Time inconsistency and free cash flow stream. *Mathematical Finance* 22, 346–378.
- Dang, D.-M. and P. Forsyth (2016). Better than pre-commitment mean-variance portfolio allocation strategies: a semi-self-financing Hamilton-Jacobi-Bellman equation approach. *European Journal of Operational Research* 250, 827–841.

- Dang, D.-M., P. Forsyth, and K. Vetzal (2017). The 4% strategy revisited: A pre-commitment optimal mean-variance approach to wealth management data. *Quantitative Finance* 17, 335–351.
- Dang, D. M. and P. A. Forsyth (2014). Continuous time mean-variance optimal portfolio allocation under jump diffusion: A numerical impulse control approach. *Numerical Methods for Partial Differential Equations* 30, 664–698.
- Dang, D. M., P. A. Forsyth, and Y. Li (2016). Convergence of the embedded mean-variance optimal points with discrete sampling. *Numerische Mathematik* 132, 271–302.
- Dichtl, H., W. Drobetz, and M. Wambach (2016). Testing rebalancing strategies for stock-bond portfolios across different asset allocations. *Applied Economics* 48, 772–788.
- Esch, D. N. and R. O. Michaud (2014). The false promise of target date funds. Working paper, New Frontier Advisors, LLC.
- Estrada, J. (2014). The glidepath illusion: an international perspective. *Journal of Portfolio Management* 40(4), 52–64.
- Forsyth, P. A. and K. R. Vetzal (2017). Robust asset allocation for long-term target-based investing. *International Journal of Theoretical and Applied Finance* 20:3. 1750017 (electronic).
- Forsyth, P. A. and K. R. Vetzal (2019). Optimal asset allocation for retirement saving: Deterministic vs. time consistent adaptive strategies. *Applied Mathematical Finance*, forthcoming.
- Gougeon, P. (2009). Shifting pensions. *Statistics Canada: Perspectives on Labour and Income* 10(5), 16–23.
- Holt, J. (2017). 2016: A growth year for target-date funds. See <https://news.morningstar.com/cover/videocenter.aspx?id=788811>.
- Kou, S. and H. Wang (2004). Option pricing under a double exponential jump diffusion model. *Management Science* 50, 1178–1192.
- Li, D. and W.-L. Ng (2000). Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance* 10, 387–406.
- Lioui, A. and P. Poncet (2016). Understanding mean variance asset allocation. *European Journal of Operational Research* 254, 320–337.
- Ma, K. and P. Forsyth (2016). Numerical solution of the Hamilton-Jacobi-Bellman formulation for continuous time mean variance asset allocation under stochastic volatility. *Journal of Computational Finance* 20:1, 1–37.
- MacDonald, B.-J., B. Jones, R. J. Morrison, R. L. Brown, and M. Hardy (2013). Research and reality: A literature review on drawing down retirement savings. *North American Actuarial Journal* 17, 181–215.
- Menoncin, F. and E. Vigna (2017). Mean-variance target based optimisation for defined contribution pension schemes in a stochastic framework. *Insurance: Mathematics and Economics* 76, 172–184.

- Patton, A., D. Politis, and H. White (2009). Correction to: Automatic block-length selection for the dependent bootstrap. *Econometric Reviews* 28, 372–375.
- Plyakha, Y., R. Uppal, and G. Vilkov (2014). Equal or value weighting? implications for asset-pricing tests. Working paper, EDHEC.
- Politis, D. and H. White (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews* 23, 53–70.
- Ruppert, P. and G. Zanella (2015). Revisiting wage, earnings, and hours profiles. *Journal of Monetary Economics* 72, 114–130.
- Tse, S. T., P. A. Forsyth, and Y. Li (2014). Preservation of scalarization optimal points in the embedding technique for continuous time mean variance optimization. *SIAM Journal on Control and Optimization* 52, 1527–1546.
- Vigna, E. (2014). On efficiency of mean-variance based portfolio selection in defined contribution pension schemes. *Quantitative Finance* 14, 237–258.
- Vigna, E. (2017). Tail optimality and preferences consistency for intertemporal optimization problems. Working paper no. 502 , Collegio Carlo Alberto, Università Degli Studi di Torino.
- Wang, J. and P. A. Forsyth (2011). Continuous time mean variance asset allocation: A time-consistent strategy. *European Journal of Operational Research* 209, 184–201.
- Zhou, X. Y. and D. Li (2000). Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization* 42, 19–33.