Towards the Design of Improved Forms of Target-Date Funds

LIONEL MARTELLINI
Professor of Finance, EDHEC Business School
Scientific Director, EDHEC Risk Institute

VINCENT MILHAU
Senior research engineer, EDHEC Risk Institute

A global trend towards the use of stricter, mark-to-market, accounting rules combined with an enhanced regulatory focus on risk management have led corporations to transfer an increasing fraction of pension-related risks to individuals. As a consequence, the retirement system in most developed countries has experienced a substantial transformation in recent years, with a shift from defined-benefit (DB) plans to defined-contribution (DC) plans, and the development of Individual Retirement Accounts. In 2009, DC assets comprise 42% of global pension assets compared with 32% in 1999 (TowersWatson [2010]), a trend that is likely to keep increasing. As a result of this trend, employees must increasingly rely on their own saving and investment decisions to fund their retirement. This is a serious concern, not only because of the induced risk transfer, but also because individual investors typically lack the expertise needed to implement educated investment decisions, and often show a great deal of inertia (see Mitchell *et al.* [2006] and Wise [2009]). In response to this concern, the asset management industry has started to design dedicated investment mutual fund products aiming at providing investors with one-stop solutions to their life-cycle investment needs. These funds, known as *target date funds* (TDFs) or *life-cycle funds*, typically rebalance the investments in the underlying funds so as to implement a deterministic decrease of equity allocations (also known as *glide path*) until a date called the *target date* or *target maturity date* of the fund. This prescription is somewhat reminiscent of the Shiller [2005] rule of thumb, which states that the allocation to equity should approximately be given by 100 minus the investor's age in years.

Embedding life-cycle investment decisions within a one-stop decision is certainly a valuable attempt at providing added-value to unsophisticated investors who otherwise will likely make sub-optimal decisions, and eventually tend to stick to them. In fact TDFs have experienced a great deal of success with total TDF assets having been multiplied more than 12-fold, to over $180bn, between 2002 and March 2008, starting with about $1 billion in 1996, when Fidelity launched its own version of these funds. The growth is accelerating, having drawn in $58bn in 2007 compared with $35bn in 2006; in the first three months of 2008, equity funds experienced outflows of more than $40bn while TDFs gathered nearly $15bn si2009. While the life cycle investing concept is unambiguously gaining traction in the industry, an outstanding question remains, however, whether one can fully rationalize the deterministic glide paths implemented in most target date fund products within the prescriptions of modern portfolio theory.

Long-term investment decisions in the presence of a stochastic opportunity set have been formally analyzed by Merton [1971], who shows that the presence of risk factors that impact the productivity of wealth justifies the introduction of intertemporal hedging demands in an investor's optimal allocation. Subsequent papers have shown that when maximizing utility from her nominal wealth, an investor only hedges those state variables that impact the nominal short-term rate and the market prices of risk (see Lioui and Poncet [2001] and Detemple *et al.* [2003] for recent references). Other papers have solved explicitly the portfolio choice problem when only one state variable is stochastic (*e.g.*, the nominal interest rate in Sørensen [1999], Lioui and Poncet [2001] and Munk and Sørensen [2004], or the stock index Sharpe ratio in kim and Omberg [1996], Campbell and Viceira [1999] or Wachter [2002]). More realistic models have also been developed to account for the presence of both state variables. This challenge has been addressed both in the context of discrete-time VAR models (see for example Campbell *et al.* [2003], and in the context of continuous-time models, either by solving numerically the Hamilton-Jacobi-Bellman (HJB) equation obtained through dynamic programming (see Brennan *et al.* [1997]), or, more recently, by exploiting the affine structure of the model (in the sense of Dai and Singleton [2000] and Liu [2007]) and solving the HJB equation explicitly (Munk *et al.* [2004]) or quasi-explicitly up to the solution of a series of ordinary differential equations (Sangvinatsos and Wachter, [2005]). A related strand of the literature has focused on portfolio choice

* lionel.martellini@edhec.edu

over the life-cycle when the investor earns a stochastic non-financial income (see for example Viceira [2001], Cocco *et al.* [2005] or Benzoni *et al.* [2007]). Overall, these papers focus on the modelization of labor income and thus assume constant investment opportunities. Munk and Sørensen (2010) relax this assumption by adding a mean-reverting short-term interest rate to a model with stochastic labor income.

In most of these papers, the fraction of wealth allocated to equities is typically shown to be a decreasing function of time-to-horizon, either because of the induced term-structure of equity risk implied by the presence of mean-reversion in equity returns, or because of the need for young investors to compensate for the bond-like nature of human capital.[1] Regardless of the explanation why young investors should hold more stocks than older investors, these findings seem to provide at best partial justification for the glide path prescription of target date funds. In fact, recommending the same allocation for investors with a given time-horizon regardless of market conditions cannot possibly be optimal in the presence of a stochastic opportunity set, and omitting such state-dependencies can lead to severe efficiency cost. As very clearly explained by Viceira and Field (2007), "(...) long-term equity investors should invest more on average in equities than their short-horizon counterparts, but they should also consider periodic revisions of this allocation as market conditions change. It is logically inconsistent to count on reduced long-term risk while ignoring the variation in returns that produces it. This market-sensitive allocation policy is very different from the asset allocation policy of life-cycle funds, whose target mix moves mechanically away from stocks as an inverse function of investment horizon, regardless of market conditions. Thus mean-reversion arguments provide, if anything, only a partial justification for the roll down schedule characteristic of life-cycle funds." In addition to Viceira and Field (2007), a number of recent studies have documented the many shortcomings of standard forms of target date funds, including Basu and Brisbane (2009), Basu and Drew (2009), Booth and Yakoubov (2000), Cairns *et al.* (2006), Lewis *et al.* (2007), or Bodie *et al.* (2009). For example, Cairns *et al.* (2006) conduct a formal analysis of the welfare loss involved in following a deterministic strategy similar to actual target date fund strategies with respect to the optimal stochastic life-cycle strategy, and they found the associated opportunity cost to be very substantial.

As a result, current forms of target date funds appear to be the wrong answer to the right question, namely how to design long-term investment strategies for financially illiterate individuals who are increasingly responsible for investment decisions related to retirement risk. The limitations of existing products extend far beyond the use of deterministic, as opposed to market-sensitive, asset allocation decisions. For example, life-cycle investment models suggest that the average decrease in equity allocation should be matched by an increased investment in the risk-free asset, as opposed to an increased investment in bonds, a feature that is not found in most existing TDFs.[2] Additionally, a large number of existing TDFs do not account for differences in risk-aversion so as to focus solely on differences in time-horizon, which is an obviously undesirable (and unneeded) simplification. Not only do current forms of TDFs can be blamed for using the wrong allocation strategies, but they should also be blamed for using the wrong building blocks. As opposed to express allocation decisions in terms of stocks versus bonds, portfolio theory unambiguously suggests that proper allocation decisions should be expressed in terms of performance-seeking portfolio versus liability-hedging portfolio. On the one hand, constant maturity bond indices, commonly used by target date fund products, are legitimate ingredients in the performance-seeking portfolio, where they will help diversify equity risk so as to generate the highest possible risk-reward ratio. On the other hand, liability-hedging portfolios should be composed of bond portfolios with a duration matching the investor's residual time-horizon. Bond indices have too long a duration for investors close to retirement and too short a duration for investors far from retirement; as a result they can hardly be regarded as safe assets for any possible investor. In addition to a poor management of interest rate risk, current TDFs also suffer from an absence of management of inflation risk, which calls for the use of inflation-linked bonds, or some other assets with attractive inflation-hedging properties, within the liability-hedging portfolio.

The intended contribution of our paper with respect to the afore-mentioned literature is two-fold. First, we contribute to the life cycle literature by proposing a comprehensive long-horizon dynamic allocation model with labor income in the presence of stochastic inflation and interest rates and a mean-reverting equity risk premium. Our model slightly extends Munk *et al.* (2004) by relaxing the assumption of a perfect negative correlation between equity returns and risk premium uncertainty.[3] We also consider real estate as an important addition to the asset mix, beyond stocks and bonds. While our setting is rich enough to account for the afore-mentioned features, as well as the presence of a (deterministic) endowment stream, we manage to obtain a quasi-explicit representation for the optimal portfolio strategy. Having quasi-analytical expressions for the optimal strategy turns out to be very useful in the analysis of the (sub)optimality of allocation strategies embedded within target date funds. We confirm that the opportunity cost involved in focusing on a deterministic allocation scheme is very substantial for reasonable parameter values. Implementing such extended forms of life-cycle investing strategies in a delegated money management context is a serious challenge, however, because it requires a narrower classification of plan participants based on factors other than the age of the participant. In a retail money management context, the challenge is to provide a parsimonious enough partition of the relevant subjective attributes (mainly age and risk-aversion) and objective attributes (in particular, current – estimated – level of the risk premium provided by equities) that can exhaust the overall population of retail investors and market conditions. Against this backdrop, our second and main contribution is to propose a formal analysis of the speed of convergence to the true optimal fully customized strategy of real-world strategies that takes into account the constraints related to limited customization in a retail money management context. Perhaps surprisingly, we find that very reasonably fine partitions, perfectly consistent with implementation in a retail money management context, allow for substantial welfare gains compared to deterministic life-cycle strategies. Our results are found to be robust with respect to the introduction of reasonable levels of measurement errors in equity risk premium, the only unobservable quantity in the model. Overall, our analysis has important potential implications for the target date fund industry since it suggests that a limited number of life-cycle investment benchmarks can be designed that will serve as relatively accurate proxies for truly optimal long-term retail investment vehicles in a retirement context.

The rest of the paper is organized as follows. In section I, we introduce a formal optimal allocation model for a long-term investor in the presence of labor income and preferences over real terminal wealth, as well as stochastic interest rates and equity risk premium. Section II proposes a numerical analysis of the model, and provides measures of welfare costs associated with following sub-optimal strategies. Section III concludes and presents suggestions for further research. Technical details and proofs of the main results are relegated to dedicated appendices.

I. Life-Cycle Investment Decisions with Stochastic Interest and Inflation Rates and a Mean-Reverting Equity Risk Premium

In this section, we solve the optimization program of an investor who faces stochastic investment opportunities, perceives an income stream, and has preferences expressed over terminal real wealth.

I.1. STATE VARIABLES AND ASSET RETURNS

We consider an investor with finite horizon date T. Uncertainty in the economy is represented through a standard probability space $(\Omega, \mathcal{A}, \mathbb{P})$ endowed with a filtration $\mathcal{F} \equiv (\mathcal{F}_t)_{t\in[0,T]}$ such that $\mathcal{F}_T = \mathcal{A}$. All processes relevant to decision making are assumed to be progressively measurable with respect to this filtration.

The sources of risk in our model are equity return uncertainty z^S, equity risk premium (Sharpe ratio) uncertainty z^λ, real estate return uncertainty z^Y, nominal interest rate uncertainty z^R, and consumer price index uncertainty z^Φ. The processes z^S, z^λ, z^Y, z^R and z^Φ are standard Wiener processes such that the vector $(z^S\, z^\lambda\, z^Y\, z^R\, z^\Phi)$ is Gaussian. We denote by ρ_{ij} the correlation between z^i and z^j, where i and j lie in the set of indices $\{R, \Phi, \lambda, S, Y\}$. Since no other source of uncertainty impacts investor's decisions, we can assume with no loss of generality that $\mathcal{F}$ is the complete version of the filtration generated by $(z^S\, z^\lambda\, z^Y\, z^R\, z^\Phi)$.

Precisely, we assume that the state variables evolve as:

$$dR_t = a(b - R_t)dt + \sigma_R dz_t^R \tag{2.1}$$

$$\frac{d\Phi_t}{\Phi_t} = \pi dt + \sigma_\Phi dz_t^\Phi \tag{2.2}$$

$$d\lambda_t^S = \kappa(\bar{\lambda} - \lambda_t^S)dt + \sigma_\lambda dz_t^\lambda \tag{2.3}$$

$$\frac{dS_t}{S_t} = (R_t + \sigma_S \lambda_t^S)dt + \sigma_S dz_t^S \tag{2.4}$$

$$\frac{dY_t}{Y_t} = (R_t + \sigma_Y \lambda_Y)dt + \sigma_Y dz_t^Y \tag{2.5}$$

or, in vector form:

$$dR_t = a(b - R_t)dt + \sigma'_R dz_t \qquad \frac{d\Phi_t}{\Phi_t} = \pi dt + \sigma'_\Phi dz_t$$

$$d\lambda_t^S = \kappa(\bar{\lambda} - \lambda_t^S)dt + \sigma'_\lambda dz_t \qquad \frac{dS_t}{S_t} = (R_t + \sigma_S \lambda_t^S)dt + \sigma'_S dz_t$$

$$\frac{dY_t}{Y_t} = (R_t + \sigma_Y \lambda_Y)dt + \sigma'_Y dz_t$$

where R denotes the nominal short-term interest rate, π is the assumed constant expected inflation rate, Φ is the price index, S is the stock index price, λ^S is its Sharpe ratio and Y is a traded real estate index price. The asset mix includes at least the stock index S, the real estate index Y and a nominal constant maturity zero-coupon bond B with a maturity denoted by τ.[4] This ensures that the risks z^S, z^Y and z^R are spanned. Hedging the risk z^λ is possible if changes in the stock value are perfectly anti-correlated to the changes in equity premium. This is equivalent to taking $\rho_{S\lambda} = -1$. Inflation risk z^Φ may be spanned by the introduction of an indexed constant maturity zero-coupon bond I. In the absence of such a bond, inflation risk is not entirely spanned, so that the market is incomplete.

Using the properties of the Vasicek model (see Vasicek [1977]), it can be shown that if τ denotes the constant maturity of the bonds, then the dynamics of B and I are given by:

$$\frac{dB_t}{B_t} = [R_t + \sigma_B(\tau)\lambda_R]dt + \sigma_B(\tau)' dz_t \tag{2.6}$$

$$\frac{dI_t}{I_t} = [R_t + \sigma_B(\tau)\lambda_R + \sigma_\Phi\lambda_\Phi]dt + \sigma_I(\tau)' dz_t \tag{2.7}$$

where $\sigma_B(\tau) = \sigma_B(\tau)\sigma_R, \sigma_I(\tau) = \sigma_\Phi + \sigma_B(\tau)\sigma_R$, and $\sigma_B(\tau) = b_a(\tau) = -\frac{1 - e^{-a\tau}}{a}$. The volatility matrix of traded assets, σ_t, is obtained by concatenating the volatility vectors of these assets. If the inflation-indexed bond is not traded, this is a 5×3 matrix. If this bond is traded, σ_t will be of size 5×4. We define λ_t as the unique market price of risk vector spanned by the assets. If the market is incomplete, *i.e.* if σ_t is not square and non-singular, there exist infinitely many prices of risk. As shown by He and Pearson (1991), the market prices of risk are those vector processes of the form $\lambda + \nu$, where ν satisfies the condition $\sigma_t' \nu_t = 0$ for all t, as well as technical measurability and integrability requirements. Each ν defines a pricing kernel $M^{\lambda+\nu}$, through:

$$dM_t^{\lambda+\nu} = M_t^{\lambda+\nu}\left[R_t\,dt + \left(\lambda_t + \nu_t\right)' dz_t\right]$$

As explained below, in this incomplete market setting, the optimization program is solved by computing the value of ν_t such that the optimal payoff is replicable by a trading strategy involving available assets.

I.2. HUMAN CAPITAL

The investor is also assumed to perceive an income flow $\left(e_t\right)_{t\in[0,T]}$, which we take to be deterministic so as to avoid further increases in the complexity of the model.[5] Let θ_t denote the vector of dollar amounts invested in the traded assets, with the following convention: the first component is the wealth allocated to the nominal bond with constant maturity τ, the second component is the amount invested in stock, the third component is the investment in real estate, and θ_t has a fourth component only if the inflation-linked bond is traded, this fourth component being the dollar investment in this indexed bond. The remaining wealth is invested in the cash, also referred to as the "risk-free asset". Thus the financial wealth evolves as:

$$dA_t = \left(A_t R_t + \theta_t' \sigma_t' \lambda_t\right)dt + \theta_t' \sigma_t' dz_t + e_t\,dt$$

Since e is a deterministic income stream, it can be valued as a bond maturing at date T and paying a continuous coupon. We define the human capital of the investor as the current value of future income streams:

$$H_t = \int_t^T B(t,s)\, e_s\, ds$$

which is a function $\mathcal{H}(t,R_t)$. In particular, the volatility vector of H is proportional to σ_R:

$$\sigma_H = \frac{\mathcal{H}_R(t,R_t)}{\mathcal{H}(t,R_t)}\sigma_R = \frac{1}{H_t}\left(\int_t^T B(t,s)\, b_a(s-t)\, e_s\, ds\right)\sigma_R$$

and the portfolio strategy θ^H replicating H is given by:

$$\theta_t^H = \frac{1}{b_a(\tau)}\left(\int_t^T B(t,s)\, b_a(s-t)\, e_s\, ds\right)e_1$$

where $e_1 = (1,0,0)'$.

It should be noted that the process $A + H$, which is the total wealth process, can be seen as the value of a self-financing strategy. Indeed, it evolves as:

$$d\left(A_t + H_t\right) = \left(A_t + H_t\right)R_t\,dt + \left[\theta_t + \theta_t^H\right]' \sigma_t' \lambda_t\, dt + \left[\theta_t + \theta_t^H\right]' \sigma_t' dz_t$$

In particular, the process $M^{\lambda}\left(A+H\right)$ follows a martingale, so that we have, for $t \le T$:

$$A_t + H_t = \mathbb{E}_t\left[\frac{M_T^{\lambda}}{M_t^{\lambda}} A_T\right] \tag{2.8}$$

I.3. OPTIMAL ALLOCATION IN THE GENERAL INCOMPLETE MARKET SETTING

Throughout the paper, we let $U(x) = \frac{x^{1-\gamma}}{1-\gamma}$ denote the investor's CRRA utility function. As in Brennan and Xia (2002), Munk *et al.* (2004) and Sangvinatsos and Wachter (2005), we assume that she has preferences over real, as opposed to nominal, wealth. Hence the problem can be mathematically written as:

$$\max_{\theta} \mathbb{E}\left[U\left(\frac{A_T}{\Phi_T}\right)\right] \tag{2.9}$$

a. Derivation of the Optimal Strategy

We first solve the problem with four assets only: the nominal bond, the stock index, the real estate index and the risk-free asset. Introducing inflation-linked bonds in the asset mix would be strongly desirable since it would allow for a perfect hedge against inflation risk. The capacity of the inflation-linked bond markets is limited, however, and we therefore focus in what follows on the general situation where perfect inflation-hedging instruments are not available.

(2.9) is a dynamic problem since the control variable is the vector of dollar amounts θ. Following the martingale approach developed by He and Pearson (1991) in incomplete market settings, this dynamic problem can be mapped into a static problem for each pricing kernel $M^{\lambda+\nu}$:

$$\max_{A_T} \mathbb{E}\left[U\left(\frac{A_T}{\Phi_T}\right)\right],\ s.t.\ \mathbb{E}\left[M_T^{\lambda+\nu} A_T\right] = A_0 + H_0 \tag{2.10}$$

The form of the budget constraint follows from (2.8). The solution technique for the portfolio choice problem then proceeds as follows: first, a candidate optimal terminal payoff is computed from (2.10) for each value of ν. Second, this payoff is priced using the pricing kernel $M^{\lambda+\nu}$. Applying Ito's lemma, one derives the diffusion term in this process and computes the value ν_t^* of ν_t such that this diffusion term is spanned by σ_S. In other words, one computes the value of ν_t such that the candidate optimal payoff is replicable by a trading strategy involving only available assets. $M^{\lambda+\nu^*}$ is called the *minimax pricing kernel*.

He and Pearson (1991) provide another useful characterization of the minimax pricing kernel (which justifies the terminology): ν^* is the value of ν that minimizes the value of the static problem (2.10) over all admissible values for ν.

Here the risky assets are the nominal bond, the stock index and the real estate index. Hence the volatility matrix of the traded assets is (we can now drop the index t because the matrix is constant over time):[6]

$$\sigma = \left(\sigma_B(\tau)\ \sigma_S\ \sigma_Y\right)$$

The spanned market price of risk vector is given by:

$$\lambda_t = \sigma\left(\sigma'\sigma\right)^{-1}\begin{pmatrix} b_a(\tau)\sigma_R\lambda_R \\ \sigma_S\lambda_t^S \\ \sigma_Y\lambda_Y \end{pmatrix}$$

Introducing the volatility matrix of traded risks:

$$\Sigma = \left(\sigma_R\ \sigma_S\ \sigma_Y\right)$$

we can rewrite λ_t as:

$$\lambda_t = \Sigma\left(\Sigma'\Sigma\right)^{-1}\begin{pmatrix} \sigma_R\lambda_R \\ \sigma_S\lambda_t^S \\ \sigma_Y\lambda_Y \end{pmatrix}$$

To isolate the effect of the stochastic λ_t^S on λ_t, we will use the following decomposition:

$$\lambda_t = \Lambda_1 + \lambda_t^S\Lambda_2$$

where:

$$\Lambda_1 = \Sigma\left(\Sigma'\Sigma\right)^{-1}\begin{pmatrix} \sigma_R\lambda_R \\ 0 \\ \sigma_Y\lambda_Y \end{pmatrix},\ \Lambda_2 = \Sigma\left(\Sigma'\Sigma\right)^{-1}\begin{pmatrix} 0 \\ \sigma_S \\ 0 \end{pmatrix}$$

We also denote by

$$N = I_5 - \sigma\left(\sigma'\sigma\right)^{-1}\sigma' = I_5 - \Sigma\left(\Sigma'\Sigma\right)^{-1}\Sigma$$

the matrix of the residual of the projection onto the columns of σ. Were the market dynamically complete, the matrix N would be zero. With these notations, we have the following result.

Proposition 1 *Consider program (2.9).*

- *The optimal payoff is given by:*

$$A_T^* = \frac{A_0 + H_0}{\mathbb{E}\left[\left(M_T^{\lambda+\nu^*}\Phi_T\right)^{1-\frac{1}{\gamma}}\right]}\left(M_T^{\lambda+\nu^*}\right)^{-\frac{1}{\gamma}}\Phi_T^{1-\frac{1}{\gamma}} \qquad (2.11)$$

where $M^{\lambda+\nu^*}$ is the minimax pricing kernel, and the optimal wealth process reads:

$$A_t^* = \frac{A_0 + H_0}{\mathbb{E}\left[\left(M_T^{\lambda+\nu^*}\Phi_T\right)^{1-\frac{1}{\gamma}}\right]}\Phi_t^{1-\frac{1}{\gamma}}\left(M_t^{\lambda+\nu^*}\right)^{-\frac{1}{\gamma}}g\left(t, R_t, \lambda_t^S\right) - H_t \qquad (2.12)$$

with:

$$g\left(t, R_t, \lambda_t^S\right) = \exp\left[\frac{1-\gamma}{\gamma}\left[\begin{array}{l} A_1(T-t) + A_2(T-t)R_t + \\ A_3(T-t)\lambda_t^S + \frac{1}{2}A_4(T-t)\left(\lambda_t^S\right)^2 \end{array}\right]\right] \qquad (2.13)$$

- *The functions A_1, A_2, A_3 and A_4 are the solutions to a system of (coupled) ordinary differential equations (ODEs):*

$$A_2(T-t) = \frac{1-e^{-a(T-t)}}{a} \qquad (2.14)$$

$$\begin{aligned} A_4'(T-t) &= \frac{\left\|\Lambda_2\right\|^2}{\gamma} + 2\left[\frac{1-\gamma}{\gamma}\sigma_\lambda'\Lambda_2 - \kappa\right]A_4(T-t) \\ &+ \frac{1-\gamma}{\gamma}\left[\sigma_\lambda^2 - (1-\gamma)\sigma_\lambda' N\sigma_\lambda\right]A_4(T-t)^2 \end{aligned} \qquad (2.15)$$

$$\begin{aligned} A_3'(T-t) &= \Lambda_2'\sigma_\Phi + \frac{\left(\Lambda_1 - \sigma_\Phi\right)'\Lambda_2}{\gamma} + \frac{1-\gamma}{\gamma}\sigma_R'\Lambda_2 A_2(T-t) \\ &+ \left[\frac{1-\gamma}{\gamma}\sigma_\lambda'\Lambda_2 - \kappa\right]A_3(T-t) \\ &+ \left[\kappa\bar{\lambda} + \frac{1-\gamma}{\gamma}\sigma_\lambda'\left(\Lambda_1 - \sigma_\Phi\right) + \frac{(1-\gamma)^2}{\gamma}\sigma_\Phi' N\sigma_\lambda\right]A_4(T-t) \\ &+ \frac{1-\gamma}{\gamma}\left[\sigma_\lambda^2 - (1-\gamma)\sigma_\lambda' N\sigma_\lambda\right]A_3(T-t)A_4(T-t) \\ &+ \frac{1-\gamma}{\gamma}\sigma_R\sigma_\lambda\rho_{R\lambda}A_2(T-t)A_4(T-t) \end{aligned} \qquad (2.16)$$

$$\begin{aligned} A_1'(T-t) &= -\pi + \frac{\left\|\Lambda_1 - \sigma_\Phi\right\|^2}{2\gamma} + \sigma_\Phi'\Lambda_1 - \frac{(1-\gamma)^2}{2\gamma}\sigma_\Phi' N\sigma_\Phi \\ &+ \left[ab + \frac{1-\gamma}{\gamma}\sigma_R'\left(\Lambda_1 - \sigma_\Phi\right)\right]A_2(T-t) + \frac{1}{2}\sigma_\lambda^2 A_4(T-t) \\ &+ \left[\kappa\bar{\lambda} + \frac{1-\gamma}{\gamma}\sigma_\lambda'\left(\Lambda_1 - \sigma_\Phi\right) + \frac{(1-\gamma)^2}{\gamma}\sigma_\Phi' N\sigma_\lambda\right]A_3(T-t) \\ &+ \frac{1-\gamma}{2\gamma}\left[\sigma_\lambda^2 - (1-\gamma)\sigma_\lambda' N\sigma_\lambda\right]A_3(T-t)^2 \\ &+ \frac{1-\gamma}{2\gamma}\sigma_R^2 A_2(T-t)^2 \\ &+ \frac{1-\gamma}{\gamma}\sigma_R\sigma_\lambda\rho_{R\lambda}A_2(T-t)A_3(T-t) \end{aligned} \qquad (2.17)$$

with the initial conditions $A_1(0) = A_3(0) = A_4(0) = 0.$

• *The optimal strategy is given by:*

$$\theta_t^* = \left(A_t^* + H_t\right)\begin{bmatrix} \frac{1}{\gamma}\left(\sigma'\sigma\right)^{-1}\sigma'\lambda_t + \left(1-\frac{1}{\gamma}\right)\frac{A_2\left(T-t\right)}{A_2\left(\tau\right)}e_1 \\ -\left(1-\frac{1}{\gamma}\right)\left[A_3\left(T-t\right)+A_4\left(T-t\right)\lambda_t^S\right] \\ \left(\sigma'\sigma\right)^{-1}\sigma'\sigma_\lambda + \left(1-\frac{1}{\gamma}\right)\left(\sigma'\sigma\right)^{-1}\sigma'\sigma_\Phi \end{bmatrix}$$

$$-\frac{1}{b_a\left(\tau\right)}\mathcal{H}_R\left(t,R_t\right)e_1 \qquad (2.18)$$

Proof. The proof of this result extends Munk *et al.* (2004) to the case where equity premium risk is not necessarily spanned. It can be found in Martellini and Milhau (2010a).

This result is completed by the following proposition, which provides an expression for the indirect utility function. This expression will be used for the computation of utility costs from following suboptimal strategies (see section 3).

Proposition 2 *The indirect utility function is:*

$$J\left(t,A_t^*,R_t,\lambda_t^S,\Phi_t\right) = \frac{1}{1-\gamma}\left(\frac{A_t^*+H_t}{\Phi_t}\right)^{1-\gamma} g\left(t,R_t,\lambda_t^S\right)^\gamma$$

b. Analysis of the Optimal Portfolio

The optimal strategy (2.18) can also be written in terms of weights as:

$$\omega_t^* = \left(1+\frac{H_t}{A_t^*}\right)\omega_t^{*,0} - \frac{1}{A_t^*}\theta_t^H \qquad (2.19)$$

where $\omega_t^{*,0}$ is the optimal weight vector when there is no income and θ_t^H is the portfolio replicating the income stream (also known as income-hedging portfolio), as defined in subsection 2.1. The portfolio strategy $\omega_t^{*,0}$ is given by:

$$\omega_t^{*,0} = \frac{1}{\gamma}\omega_t^{PSP} - \left(1-\frac{1}{\gamma}\right)\left[A_3\left(T-t\right)+A_4\left(T-t\right)\lambda_t^S\right]$$
$$\omega^\lambda - \left(1-\frac{1}{\gamma}\right)A_2\left(T-t\right)\omega^R + \left(1-\frac{1}{\gamma}\right)\omega^\Phi$$

where:

• $\omega_t^{PSP} = \left(\sigma'\sigma\right)^{-1}\sigma'\lambda_t$ is the performance-seeking portfolio (PSP);

• $\omega^R = \left(\sigma'\sigma\right)^{-1}\sigma'\sigma_R$ is the interest rate-hedging portfolio;

• $\omega^\lambda = \left(\sigma'\sigma\right)^{-1}\sigma'\sigma_\lambda$ is the equity premium-hedging portfolio;

• $\omega^\Phi = \left(\sigma'\sigma\right)^{-1}\sigma'\sigma_\Phi$ is the inflation-hedging portfolio.

Another way of decomposing the optimal allocation for the individual investor would consist in grouping the intertemporal hedging demand against interest rate risk ω^R and the intertemporal hedging demand against inflation risk ω^Φ so as to identify a portfolio that would serve as a proxy for the safe asset for a long-term investor preparing for retirement, a portfolio that would typically be referred to as a *liability-hedging portfolio* in a pension fund context. In a complete market setting, this portfolio would be 100% invested in a pure discount inflation-linked bond with a maturity date matching the investor's horizon. In the absence of inflation-linked bonds, the liability-hedging portfolio can be written as $\omega^\Phi - A_2(T-t)\,\omega^R$. On the one hand, the term $-A_2(T-t)\,\omega^R$ is the (perfect) hedge against interest rate risk, which can be interpreted as the replicating portfolio for a pure discount bond with maturity T using the constant maturity bond index and cash. On the other hand, ω^Φ is the imperfect hedge against inflation risk, which can be interpreted as the portfolio of risky assets that exhibits the maximum correlation with respect to changes in inflation.

A number of additional comments are in order. It should first be noted that $\omega_t^{*,0}$ is independent of the current wealth of the investor, a property which is typical of CRRA utility functions. We also confirm that the optimal strategy involves not only time-dependencies, but also state-dependencies, in contrast to the heuristic deterministic glide paths of life cycle funds. In particular, the optimal allocation will change as a function of changes in the stock index Sharpe ratio. Given that the functions A_i are continuous at zero, the hedging demand ω_t^λ against the Sharpe ratio risk vanishes when the time-to-maturity shrinks to zero, as expected. For a finite time-to-maturity, this demand is non zero. In the specific case of a perfect anti-correlation between shocks to S and shocks to λ^S, this hedging demand simplifies into a full invested in the stock index. In the general case $\rho_{S\lambda} > -1$, some general properties can be analyzed even in the absence of fully analytical expressions.[7] In particular, it can be shown that for $\gamma > 1$, A_4 is a positive increasing function. This result implies that if the investor is more risk-averse than the logarithmic one (*i.e.*, if $\gamma > 1$), then an increase in the expected stock return should lead to a decrease in the allocation to the portfolio hedging against Sharpe ratio risk. In particular, if $\rho_{S\lambda} > -1$, an increase in the Sharpe ratio should result in a higher allocation to stocks. It should be noted that this effect is independent from the effect of a higher Sharpe ratio on the composition of the tangency portfolio.[8]

Secondly, for a vanishing time-to-maturity, the human capital and the hedging demand against income risk in (2.18) goes to zero, so that the investor behaves as if she were to receive no income:

$$\lim_{T-t\downarrow 0}\left[\omega_t^* - \omega_t^{*,0}\right] = 0$$

Moreover, the hedging demands against interest rate risk and equity premium risk cancel as well. Eventually, we obtain that:

$$\lim_{T-t\downarrow 0}\omega_t^* = \frac{1}{\gamma}\omega_t^{PSP} + \left(1-\frac{1}{\gamma}\right)\omega^\Phi$$

If there were no inflation risk, the investor would behave myopically, by investing only in the PSP and the cash. If inflation is stochastic, the investor keeps hedging this risk, even just before the horizon.

Thirdly, the decomposition (2.19) sheds light on the impact of the non-financial income. It shows that in the presence of deterministic income, the investor chooses the same portfolio as an other investor with the same age and risk aversion characteristics who would receive no income but would have a wealth equal to $A^* + H$. This result holds more generally in the presence of stochastic income as long as income risk is spanned (see Munk and Sørensen (2010)).

We now turn to a numerical analysis of the model, in which we shall make the following assumptions. First, we take the income stream to be zero. Indeed, as explained above, an investor with non-zero income behaves as an investor without income and with a larger wealth. Second, we take the stock index Sharpe ratio to be perfectly anti-correlated with stock index returns. This assumption is supported by empirical works that proxy the expected excess return as an affine function of the dividend yield (see *e.g.* Barberis (2000) and Xia (2001), who find a strongly negative correlation between excess returns and the dividend yield), and has been also made by Wachter (2002) and Munk *et al.* (2004). It is not crucial but it clarifies the impact of the stochastic Sharpe ratio on the allocation: the hedging portfolio ω^λ is entirely invested in equities, so that an increase in the Sharpe ratio leads to a higher allocation to the stock index, both in the speculative portfolio and in the hedging demand. Although equity premium risk is spanned by the stock itself, the market is still incomplete because we do not necessarily assume that a perfect inflation-hedging instrument exists.

II. Numerical Analysis of the Model

The goal of this section is to provide a quantitative estimate for the utility loss incurred by following a heuristic, sub-optimal, strategy such as the one typically implemented by currently available target date funds. To assess formally the utility loss incurred by following a sub-optimal strategy with respect to the continuous-time optimal strategy, we compute the associated Monetary Utility Loss (MUL), defined as the amount x, expressed as a percentage of initial wealth, such that the investor is indifferent in terms of expected utility between the following two options: (a) follow the sub-optimal strategy starting with the initial capital A_0; (b) follow the optimal (continuous-time) strategy starting with a lower initial capital $A_0 - x$. In the subsequent numerical exercise, all real-world strategies we test strategies are implemented using a quarterly time-step, and the expected utility is approximated using 50,000 simulated outcomes for the terminal payoff. This will allow to provide a quantitative measure for the welfare loss for an investor involved in following a sub-optimal strategy based on a deterministic glide path with respect to the idealized continuous-time and fully customized stochastic life-cycle strategy, as in Cairns *et al.* (2006), for example. In addition to estimating monetary utility loss, we also estimate for various optimal or heuristic strategies a *return-to-risk* (or *risk-return*) ratio, which is the ratio of the mean to the standard deviation of the annualized real log return. Formally, this ratio is defined as:

$$\text{Risk-return ratio} = \frac{\mathbb{E}[r_T]}{\sqrt{\mathbb{V}[r_T]}}$$

where $r_T = \frac{1}{T}\ln\left(\frac{A_T}{\Phi_T}\frac{I(0,T)}{A_0}\right)$ is the log return on real wealth. Obviously, outperformance as measured by this return-to-risk ratio is not strictly equivalent to outperformance as measured in terms of expected utility of terminal wealth, which is the maximization objective used to obtain the optimal strategy. It is, however, a reasonable indicator of quality for a given portfolio strategy from a mean-variance perspective.

Our base case parameters for the stock price, the Sharpe ratio and the price index processes are borrowed from Munk *et al.* (2004). In particular, we set the correlation between the Brownian motions z^S and z^λ to -1, which removes the incompleteness due to the Sharpe ratio. The maturity of the constant-maturity bond is taken equal to 10 years.[9] Parameters for the real estate process are taken from De Jong *et al.* (2007). All parameter values are displayed in table 1.[10] We also report the nominal yields for three different horizons (1 year, 10 years and 20 years), as well as the annualized expected rates of return on stocks and on the price index. That the annualized expected inflation rate is constant across horizons follows from the Geometric Brownian motion assumption made on the price index. It is, however, less clear why the annualized expected rate of return on stocks is independent of the investment horizon. In fact, it can be shown that this property is only obtained when the initial values of Sharpe ratio, λ_0^S and short-term rate, R_0, are equal to their long-term means, which is assumed to hold true in our base case.

It should be recognized at this stage that the continuous-time optimal strategy is a purely idealized strategy that could not be implemented in realistic situations, and as such proves to be an unfair benchmark for the heuristic deterministic glide path strategies. In fact real-world implementations of stochastic life-cycle strategies in a retail money management mutual fund context will suffer from at least three limitations with respect to the idealized optimal strategy. Firstly, the strategy will involve a discrete partition of the set of trading dates: the strategy will be implemented with discrete, as opposed to continuous trading, with typically a monthly or quarterly trading period. The discrete implementation of the optimal strategy will involve a monetary utility loss, which needs to be measured in the first place before assessing the inefficiency of heuristic target date fund strategies. Secondly, the strategy will also involve a discrete partition of the set of investors. Indeed, in a retail money management context, the strategy will not be fully customized to fit each single retail investor's profile, and investors will instead have to be sorted in various categories as a function of their age (or time-horizon), in such a way that in

Table 1. Base case parameters

Panel A displays the base case parameter values. The dynamics of the state variables are given in equations (2.1), (2.2), (2.3), (2.4) and (2.5). Panel B shows the nominal zero-coupon yields implied by these values, as well as the annualized expected logarithmic rates of return on the stock and on the price index.

Panel A: Continuous-time parameter values	
Stock index	
σ_S	0.1468
S_0	100
Sharpe ratio	
σ_λ	0.047
κ	0.0608
$\bar{\lambda}$	0.4414
λ_0	0.4414
Nominal rate	
σ_R	0.0195
λ_R	– 0.2747
a	0.0395
b	0.0369
R_0	0.0369
Price index	
σ_Φ	0.0081
π	0.0357
Φ_0	1
Real estate	
σ_Y	0.596
λ_Y	0.734
Y_0	100
Correlations	
ρ_{RS}	0.0845
$\rho_{S\Phi}$	– 0.0678
$\rho_{R\Phi}$	0.0032
$\rho_{R\lambda}$	– 0.0845
$\rho_{\Phi\lambda}$	0.0678
$\rho_{S\lambda}$	– 1
ρ_{SY}	– 0.0134
ρ_{RY}	– 0.0036
$\rho_{\Phi Y}$	0.088
$\rho_{Y\lambda}$	0.0134
Panel B: Implied values	
Nominal yields (in %)	
1 year	3.95
10 years	5.57
20 years	6.42
Annualized expected stock returns (in %)	
Any horizon	9.09
Annualized expected inflation (in %)	
Any horizon	3.57

general to an investor with a given time-horizon will be assigned a strategy that will only be approximately optimal. For example, an investor with a time-horizon equal to say 19 years or 21 years will be placed at the initial date within a category defined by a 20 year time-horizon, and the strategy that will be proposed to all investors in that particular category will only be optimal for investors having a time-horizon *exactly* identical to 20 years.[11] Thirdly, the strategy will involve a discrete partition of the set of market conditions: as opposed to using the exact risk premium estimate at each given point in time and for each state of the world, risk premium levels will instead be typically sorted in various categories such as low, average and high risk premium values.

As a result, beyond measuring the welfare cost of strategies based on deterministic glide paths with respect to theoretical idealized optimal continuous-time and continuous-state fully customized strategies (an unfair competition because such idealized strategies will never be implemented in practice), our main focus and contribution is to compare the heuristic strategies based on deterministic glide paths with respect to real world versions of the optimal strategies involving a discrete partition of the sets of trading dates, investors' characteristics and market conditions. After all, if it turns out that the reported inefficiency of deterministic life-cycle strategies happens to be substantially lower when compared to real-world versions of the stochastic life-cycle strategies, then any hope for generating a new improved version of target date funds will have to be re-assessed.

II.1. INTRODUCING A DISCRETE PARTITION OF THE SET OF TRADING DATES

To provide a fairer benchmark to heuristic strategies, we first recognize that optimal stochastic life-cycle strategies would not be implemented in continuous-time. We therefore consider a discrete version of the optimal strategy, which we call "Discretized Optimal Strategy", and where the vector of weights at time t_i is taken to be $\omega^*_{t_i}$. In terms of the heuristic deterministic life-cycle strategy that gradually switches from stocks and real estate ("risky assets") to the bond index (the "safe" asset), we make the following assumption. Formally, the "risky" part of the allocation is taken to be a portfolio of 80% stocks and 20% real estate, and the "safe" part is fully invested in the bond index:

$$\omega^{\text{risky}} = (0, 0.8, 0.2)', \quad \omega^{\text{safe}} = (1, 0, 0)'$$

The weight vector of the deterministic life-cycle strategy is given by:

$$\omega^{dlf}_{t_i} = f(t_i)\omega^{\text{risky}} + \left[1 - f(t_i)\right]\omega^{\text{safe}}$$

where $f(t)$ is a deterministic function of time, decreasing from 80% at the initial date to 20% just before the horizon, by annual steps. Mathematically, if time is expressed in years, we have:

$$f(t_i) = 0.8 - \frac{0.6}{T-1}\lfloor t_i \rfloor$$

For example, in the case of a 20-year horizon, the allocation to the risky part is decreased by $0.6/19 \approx 316$ basis points per year. In unreported results, we have also tested other variants of deterministic schemes similar to those implemented in practice, and have obtained very similar results in terms of the induced welfare loss.

Figure 1 provides a measure of welfare loss for this deterministic life-cycle strategy with respect to the optimal idealized continuous-time strategy for three values of the risk aversion parameter γ and an initial time-horizon taken to be 20 years (throughout the paper, we maintain the assumption of an initial time-horizon of 20 years; in unreported results, we have tested time-horizon values of 10 years and 30 years, and have obtained qualitatively similar results). The welfare loss appears to very substantial, around 80% for all tested risk-aversion levels. By contrast, the introduction of quarterly rebalancing constraints only implies a very modest welfare loss with respect to a continuous-time implementation of the optimal strategy, equal to 6.0% for $\gamma = 10$, 4.2% for $\gamma = 15$, and 2.9% for $\gamma = 20$. Overall, these results suggest that ample room is available for improvements of the heuristic target date funds currently available. Before turning to an assessment of the impact of the introduction of a realistic discrete partition of the set of investors and market conditions, we also consider a variant of the optimal strategy implemented in discrete time, where the real estate index is excluded from the menu of traded assets. Obviously, ignoring the real estate index implies a substantial welfare loss, reaching 43.3% for $\gamma = 10$, 31.9% for $\gamma = 15$, and 25.5% for $\gamma = 20$. These results confirm that an investable real estate index can be a useful addition to the menu of asset classes in the context of long-term investment strategies. In fact, real estate can in general be usefully introduced in the performance-seeking portfolio where it allows for diversification benefits with respect to stocks and bonds, as well as in the inflation-hedging portfolio.[12]

To give a better sense of the performance of the deterministic and of the optimal discretized strategy, we report in table 2 the annualized expected (logarithmic) returns on these strategies, as well as their annualized volatilities. As the investment horizon becomes more distant, the annualized expected return is slightly decreasing with the horizon, but the annualized volatilities of discretized optimal strategies exhibit a strongly decreasing pattern.

II.2. INTRODUCING A DISCRETE PARTITION OF THE SET OF INVESTORS AND MARKET CONDITIONS

We now explain how the real-world implementation constraints can be integrated in the simulation of the various life-cycle investing strategies. In particular, we describe how the partition is performed in terms of objective investment attributes, *i.e.*, market conditions such as summarized in terms of different values of the Sharpe ratio, and subjective investment attributes, in this case time-horizon.

We first describe the partition of the set of market conditions as a function of the stock index Sharpe ratio, which is assumed to follow the mean-reverting process (2.3). As is well-known from the properties of Ornstein-Uhlenbeck processes, the long-term mean and variance are respectively $\bar{\lambda}$ and $\sigma_\lambda^2/(2\kappa)$. The crudest possible partition of market conditions would involve replacing all realizations of the process λ^S by the constant long-term mean $\bar{\lambda}$. A somewhat finer partition of the set of market conditions would involve distinguishing between high, moderate and low risk premium levels. In what follows, we use the following threshold values:

$$\lambda_{\text{inf}} = \bar{\lambda} - 2 \times 1.96\sqrt{\frac{\sigma_\lambda^2}{2\kappa}},\ \bar{\lambda}, \qquad (3.1)$$
$$\lambda_{\text{sup}} = \bar{\lambda} + 2 \times 1.96\sqrt{\frac{\sigma_\lambda^2}{2\kappa}}$$

Table 2. Annualized expected returns and volatilities(in%)

This table reports annualized expected logarithmic rates of return on the deterministic life-cycle strategy as well as on the optimal discretized strategy for different values of the risk aversion parameter γ. It also reports the annualized standard deviations on these strategies. Parameters are fixed at the values in table 1.

		Horizon (years)		
Strategies		**1**	**10**	**20**
Deterministic life-cycle	Mean	15.21	13.86	12.35
	Volatility	14.63	13.06	15.72
Discretized optimal (γ = 10)	Mean	16.67	16.19	15.08
	Volatility	27.27	16.63	8.98
Discretized optimal (γ = 15)	Mean	13.91	13.45	12.40
	Volatility	26.15	14.97	6.09
Discretized optimal (γ = 20)	Mean	12.43	12.00	10.99
	Volatility	25.77	14.33	4.63

Figure 1. Monetary utility losses from following heuristic strategies

These figures represent monetary utility losses from following various heuristic strategies. All utility losses are computed with respect to the indirect utility from implementing the optimal strategy in continuous time. "Disc. Opt." refers to the optimal strategy implemented on a discretized basis; "Det. LC" corresponds to a deterministic life-cycle strategy based on a "rule of thumb" close to current practice; "NoRE" refers to a variant of the optimal strategy implemented in discrete time, where the real estate index is excluded from the menu of traded assets. Unless otherwise indicated, parameters are fixed at their base case values (see table 1).

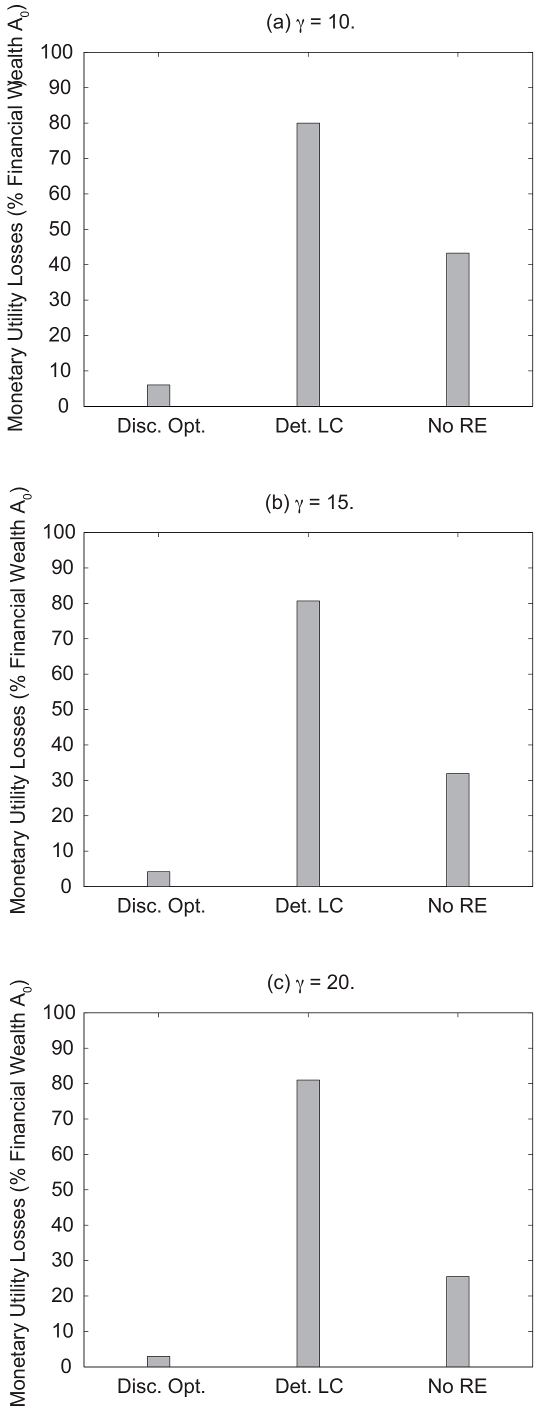

The classification rule is that the realization of λ^S at a given date and for a given state of world will be replaced by the closest of these threshold three values (if λ^S is equally spaced from two of these values, the lower one is considered). Therefore, the second level of partition contains three classes, which are:

$$\left]-\infty, \frac{\lambda_{\inf} + \bar{\lambda}}{2}\right], \left]\frac{\lambda_{\inf} + \bar{\lambda}}{2}, \frac{\bar{\lambda} + \lambda_{\sup}}{2}\right], \left]\frac{\bar{\lambda} + \lambda_{\sup}}{2}, \infty\right[$$

and three standard values: $\lambda_{\inf}$, $\bar{\lambda}$ and $\lambda_{\sup}$. The coefficient 2×1.96 in (3.1) has been chosen to ensure that the central class will contain 95% of the long-term observations of the Sharpe ratio. Values outside this interval are considered to be outliers, so that refining the partition outside the central range would be irrelevant. Therefore, one subsequently refines the partition by dividing the central interval into two equal sub-intervals, a process that will eventually converge to a continuous set of partitions. For example, the third level of partition involves 5 threshold values for the Sharpe ratio:

$$\lambda_{\inf}, \frac{\lambda_{\inf} + \bar{\lambda}}{2}, \bar{\lambda}, \frac{\bar{\lambda} + \lambda_{\sup}}{2}, \lambda_{\sup}$$

and 5 related classes. More generally, it can be verified that for $k \geq 2$, the partition of index k has $2^{k-1} + 1$ standard values. Let us denote these values by $\lambda_j^{(k)}$, for $1 \leq j \leq 2^{k-1} + 1$. In particular, the values $\lambda_j^{(2)}$ for $1 \leq j \leq 3$ are given by (3.1). Mathematically, the values for this partition can be computed as:

$$\lambda_j^{(k)} = \frac{1}{2^{k-2}}\left[\left(2^{k-2} - r\right)\lambda_{i+1}^{(2)} + r\lambda_{i+2}^{(2)}\right], \quad (3.2)$$
$$1 \leq j \leq 2^{k-1}$$

with the unique decomposition $j - 1 = 2^{k-2} i + r$ and $0 \leq r \leq 2^{k-2} - 1$. The last value is:

$$\lambda_{2^{k-1}+1}^{(k)} = \lambda_3^{(2)}$$

In terms of partitioning of the set of investors with respect to the time-horizon $T - 1$, the imperfect customization inherent to the retail money management context implies replacing the actual time-to-horizon with some approximate value. In what follows, we will consider investors with time-horizons ranging from 0 to 30 years. As a result, the crudest possible partition would involve using the same average value $\bar{T} = 15$ years for all investors, regardless of their actual time-horizon. It should be noted at this stage that all investors, including the one who happens to have a time-horizon originally equal to 15 years, will incur a welfare loss as a result of this imperfect partition. This is because for a given investor, say the one with an initial time-horizon of 15 years, residual time-to-horizon will decrease as time goes by, while the strategy will be kept implemented with a constant 15 year horizon. A slightly finer partition would involve defining three possible time-horizons:

$$T_{\inf} = 0,\ \bar{T},\ T_{\sup} = 30$$

More generally, the partition of index k would involve $2^{k-1} + 1$ threshold values, very similarly to what has been explained above for the partition of Sharpe ratio levels. In the limit of k going to infinity, a set of different portfolio strategies corresponding to a continuum of time-horizons will be offered to investors so that at all points in time the investment strategy that will be proposed to them will be perfectly consistent with their respective residual time-horizon. For finite k values, on the other hand, some utility loss will be incurred, and our motivation is to analyze the speed of convergence to zero of the welfare loss as a function of the fineness/coarseness of the partition measured by k.

Because risk-aversion is also an important parameter that should affect the optimal allocation strategy, we present three different life cycle strategy benchmarks for three different risk-aversion levels. As explained above, we do not analyze the impact of a discrete partition of the set of investors in terms of risk-aversion; because this parameter is not observable, investors are assumed to be perfectly represented by a self-selected risk-aversion profile.

a. Impact of a Discrete Partition of the Set of Investors

Figure 2 analyzes the impact of partitioning the set of time-to-horizons. It displays the MUL for the strategy based on the partition as a function of the index k. As could be expected, this MUL converges to the MUL of the discretized optimal strategy as k grows to infinity. Perhaps surprisingly, we find that the speed of convergence is very fast. In fact, for $k = 4$, which corresponds to nine classes, the MUL for the partitioned strategy already almost coincides with the MUL for the fully customized strategy implemented in discrete-time. Even assuming a constant time-horizon of 15 years for all investors at all dates leads to a welfare loss that is still lower than the welfare loss incurred with heuristic deterministic life-cycle investing strategies. In figure 3, we confirm that return-to-risk ratios are approximately twice as good for the optimal discretized strategy compared to the heuristic deterministic strategies, and the speed of convergence is again very fast. For $k = 3$, which corresponds to five classes, the return-to-risk ratio obtained with the optimal strategy implemented with a discrete partition of the set of time-horizon is already very close to what would be achieved by the fully customized version. Overall these results seem to confirm that currently available forms of TDFs are severely inefficient investment strategies, which could be improved upon in realistic environments integrating the retail money management constraints.

b. Impact of a Discrete Partition of the Set of Market Conditions

We now turn to an analysis of the impact of partitioning the set of Sharpe ratios. This is the purpose of figure 4, which reports the MUL for strategies using an approximate value for the Sharpe ratio. Again, the utility loss is decreasing with the index of the partition, and it rapidly

Figure 2. Monetary utility losses from partitioning the set of time-to-horizons

The solid line represents monetary utility losses (MUL) from following strategies based on partitions of the set of time-to-horizons. The dash-dot line represents the MUL from implementing the optimal strategy in discrete time, and the dashed line is the MUL from implementing a deterministic life-cycle strategy. Unless otherwise indicated, parameters are set at their base case values (see table 1). k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

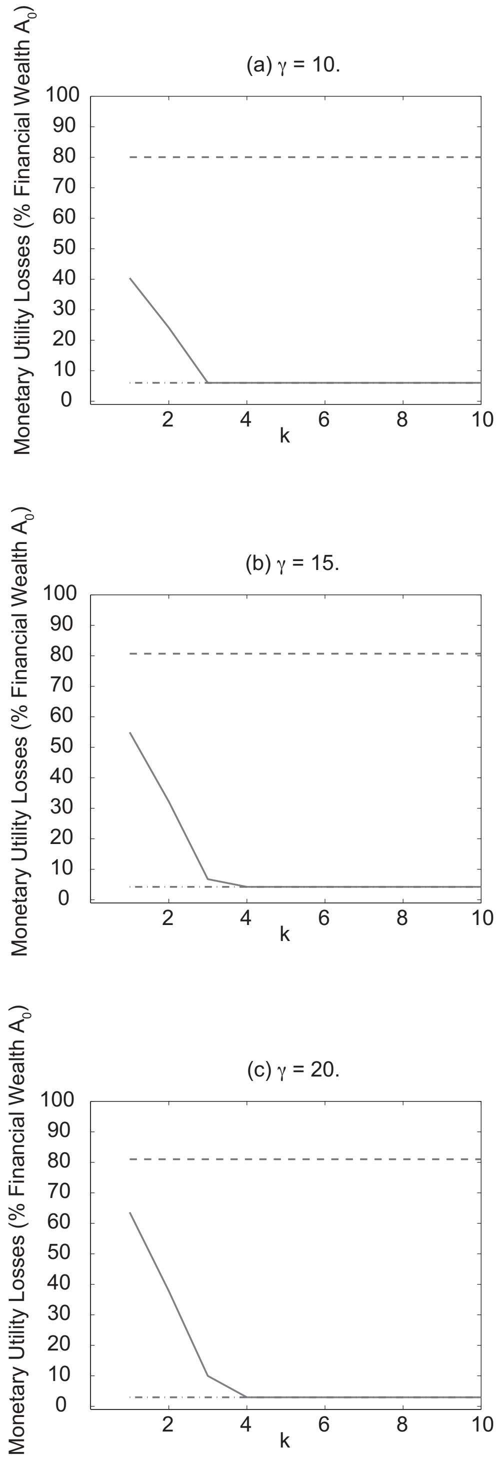

Figure 3. Risk-return ratios for strategies based on a partition of the set of time-to-horizons

The solid lines represent strategies based on partitions of the set of time-to-horizons. The dash-dot lines represent the optimal strategy implemented in discrete time, and the dashed lines represent a deterministic life-cycle strategy. The risk-return ratio is annualized. Unless otherwise indicated, parameters are set at their base case values (see table 1). k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

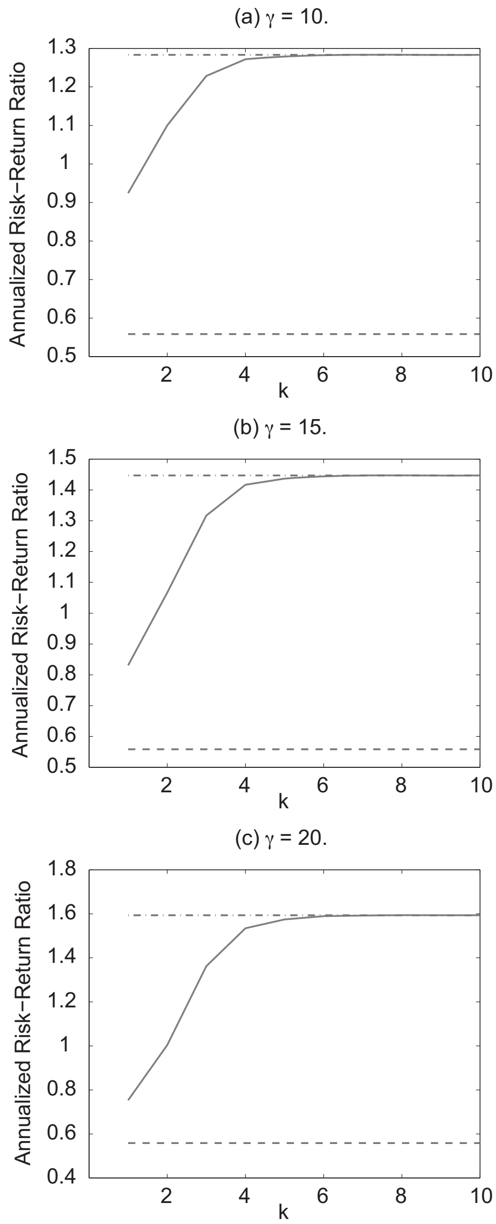

Figure 4. Monetary utility losses from partitioning the set of Sharpe ratios – Low variability

The solid line represents monetary utility losses (MUL) from following strategies based on partitions of the set of Sharpe ratios. The dash-dot line represents the MUL from implementing the optimal strategy in discrete time, and the dashed line is the MUL from implementing a deterministic life-cycle strategy. Unless otherwise indicated, parameters are set at their base case values (see table 1). k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

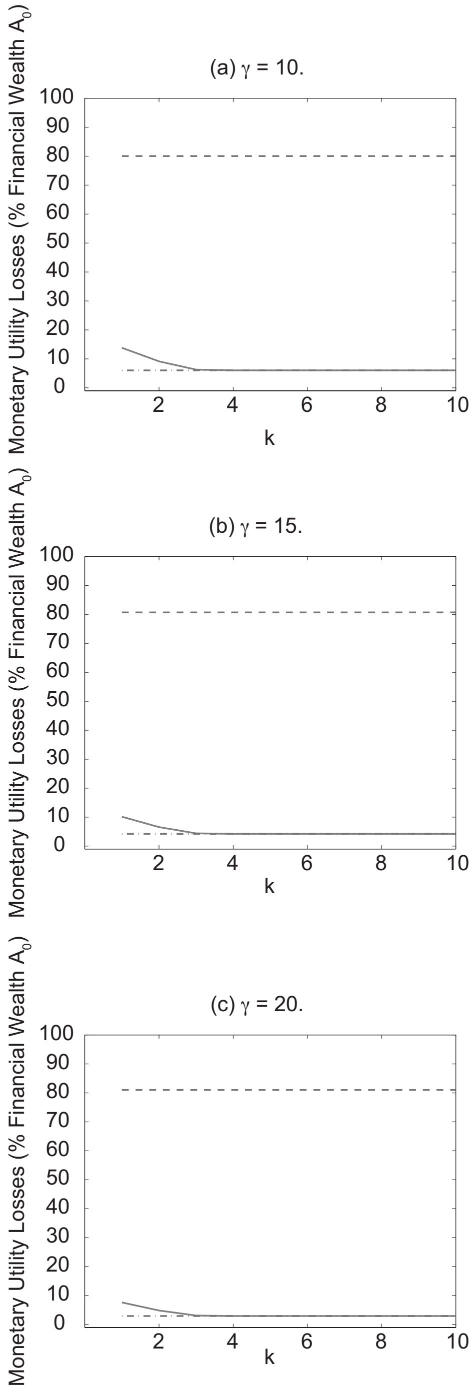

Figure 5. Risk-return ratios for strategies based on a partition of the set of Sharpe ratios – Low variability

The solid lines represent strategies based on partitions of the set of Sharpe ratios. The dash-dot lines represent the optimal strategy implemented in discrete time, and the dashed lines represent a deterministic life-cycle strategy. The risk-return ratio is annualized. Unless otherwise indicated, parameters are set at their base case values (see table 1). k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

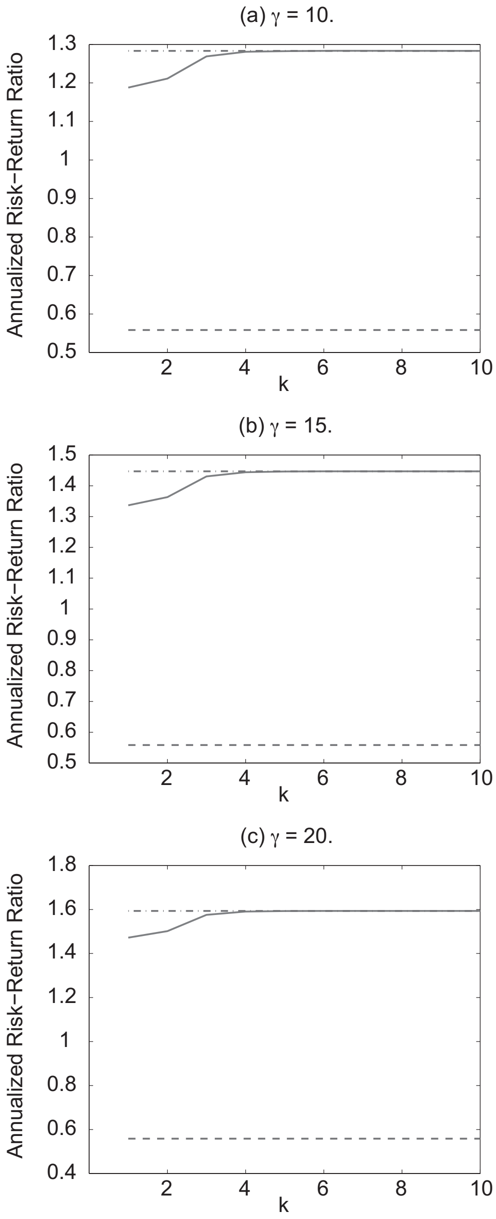

Figure 6. Monetary utility losses from partitioning the set of Sharpe ratios – Low variability

The solid line represents monetary utility losses (MUL) from following strategies based on partitions of the set of Sharpe ratios. The dash-dot line represents the MUL from implementing the optimal strategy in discrete time, and the dashed line is the MUL from implementing a deterministic life-cycle strategy. Unless otherwise indicated, parameters a reset at their base case values (see table 1), except for the standard deviation σ_λ, which was taken to be 9.4%. k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

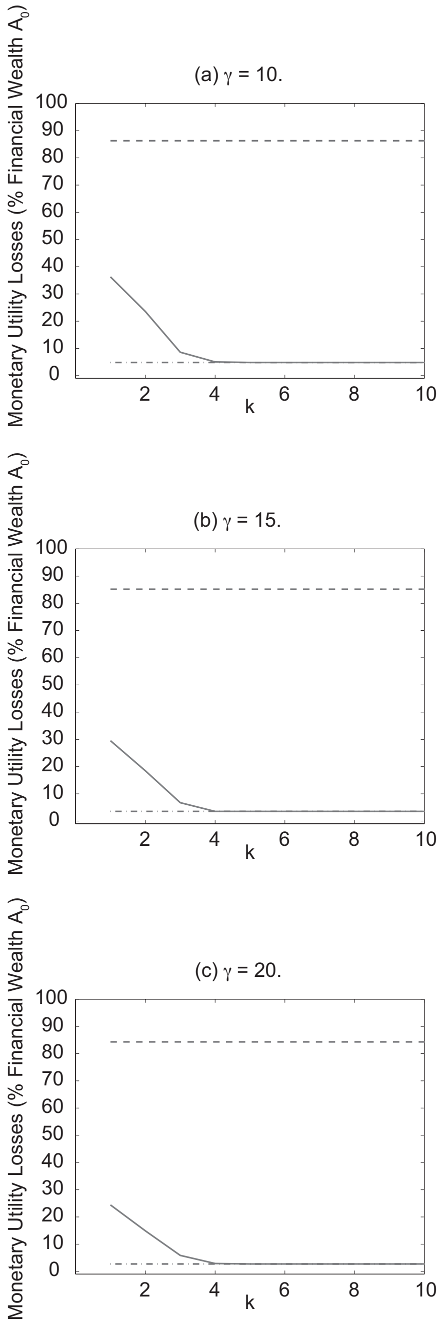

Figure 7. Risk-return ratios for strategies based on a partition of the set of Sharpe ratios – Higher variability

The solid line represents the annualized risk-return ratio over T years for strategies based on partitions of the set of Sharpe ratios. The dash-dot line represents the optimal strategy implemented indiscrete time, and the dashed line represents a deterministic life-cycle strategy. Unless otherwise indicated, parameters are set at their base case values (see table 1), except for the standard deviation σ_λ, which was taken to be 9.4%. k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

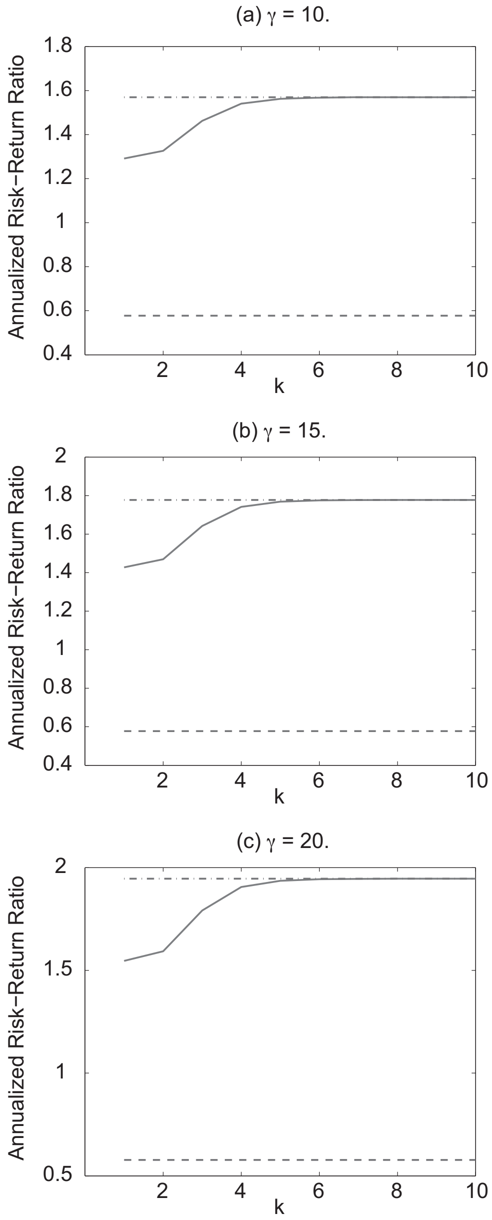

Figure 8. Monetary utility losses from partitioning both the set of time-to-horizons and the set of Sharpe ratios

The solid line represent monetary utility losses (MUL) from following strategies based on partitions of the set of Sharpe ratios and the set of horizons. The solid line represents strategies based on such partitions. The dash-dot line represents the MUL from implementing the optimal strategy indiscrete time, and the dashed line is the MUL from implementing a deterministic life-cycle strategy. Unless otherwise indicated, parameters are set at their base case values (see table 1). k denotes the index of the partition of the set of Sharpe ratios, so that the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

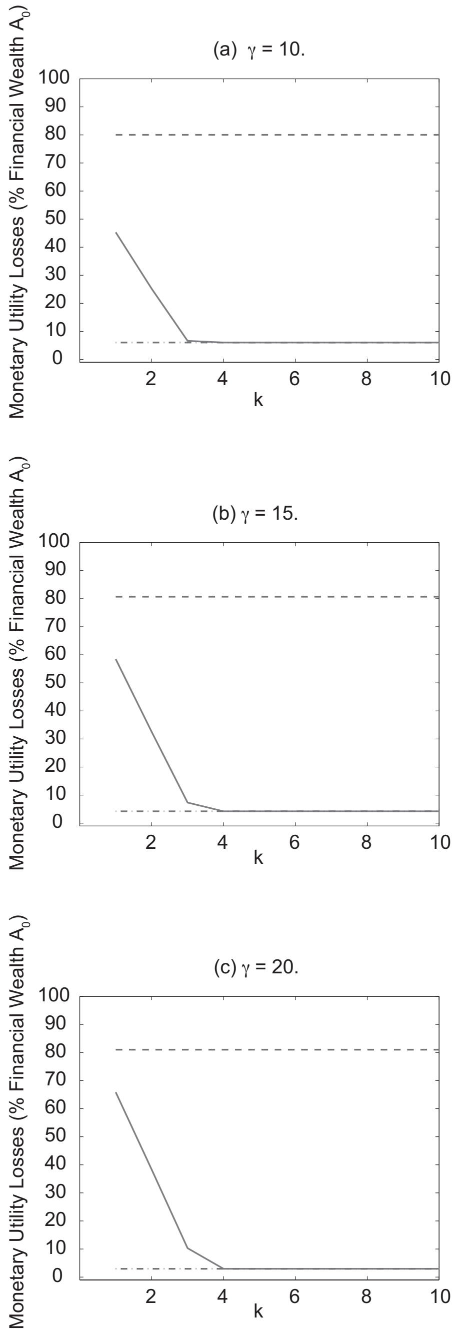

Figure 9. Risk-return ratios for strategies based on partitions of the set of time-to-horizons and of the set of Sharpe ratios

The solid lines represent strategies based on partitions of the set of time-to-horizons and of the set of Sharpe ratios. The dash-dot line represents the optimal strategy implemented indiscrete time, and the dashed line represents a deterministic life-cycle strategy. The risk-return ratio is annualized. Unless otherwise indicated, parameters are set at their base case values (see table 1). k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

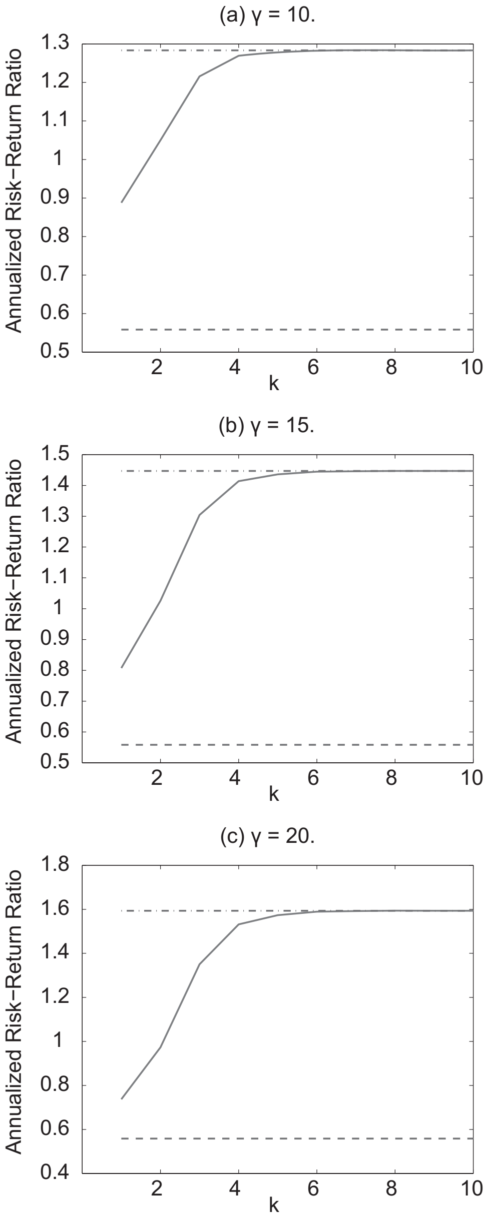

Figure 10. Monetary utility losses for strategies based on partitions of the set of time-to-horizons and of the set of Sharpe ratios when the Sharpe ratio is imperfectly observed – Low volatility of the noise process

The solid lines represent strategies based on partitions of the set of time-to-horizons and of the set of Sharpe ratios. The dash-dot line represents the optimal strategy implemented in discrete time, and the dashed line represents a deterministic life-cycle strategy. Unless otherwise indicated, parameters are set at their base case values (see table 1), and the Sharpe ratio is observed with a noise whose standard deviation is twice the long-term standard deviation of the stochastic λ^S. k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1} + 1$. γ is the relative risk aversion coefficient of the investor.

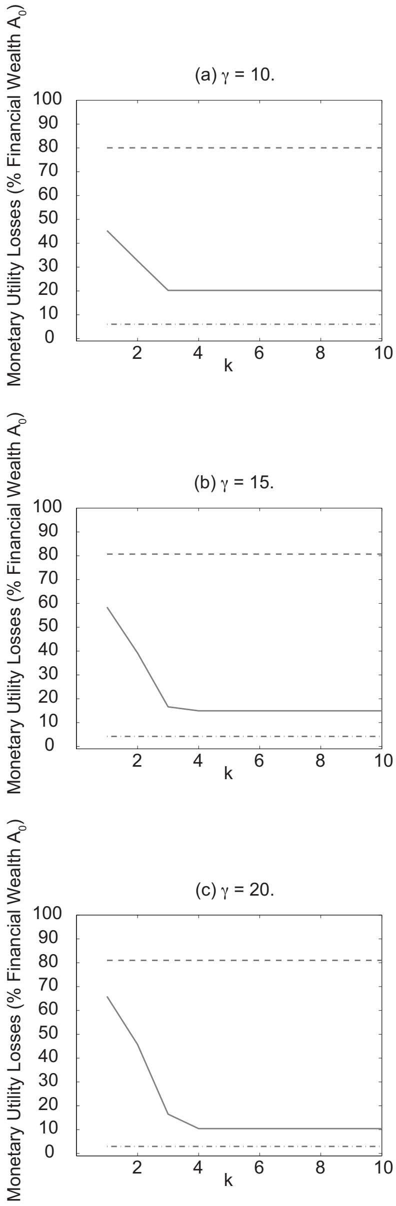

Figure 11. Monetary utility losses for strategies based on partitions of the set of time-to-horizons and of the set of Sharpe ratios when the Sharpe ratio is imperfectly observed – High volatility of the noise process

The solid lines represent strategies based on partitions of the set of time-to-horizons and of the set of Sharpe ratios. The dash-dot line represents the optimal strategy implemented in discrete time, and the dashed line represents a deterministic life-cycle strategy. Unless otherwise indicated, parameters are set at their base case values (see table 1), and the Sharpe ratio is observed with a noise whose standard deviation is five times the long-term standard deviation of the stochastic λ^S. k denotes the index of the partition, so that for $k \geq 2$, the number of classes in the partition is $2^{k-1}+1$. γ is the relative risk aversion coefficient of the investor.

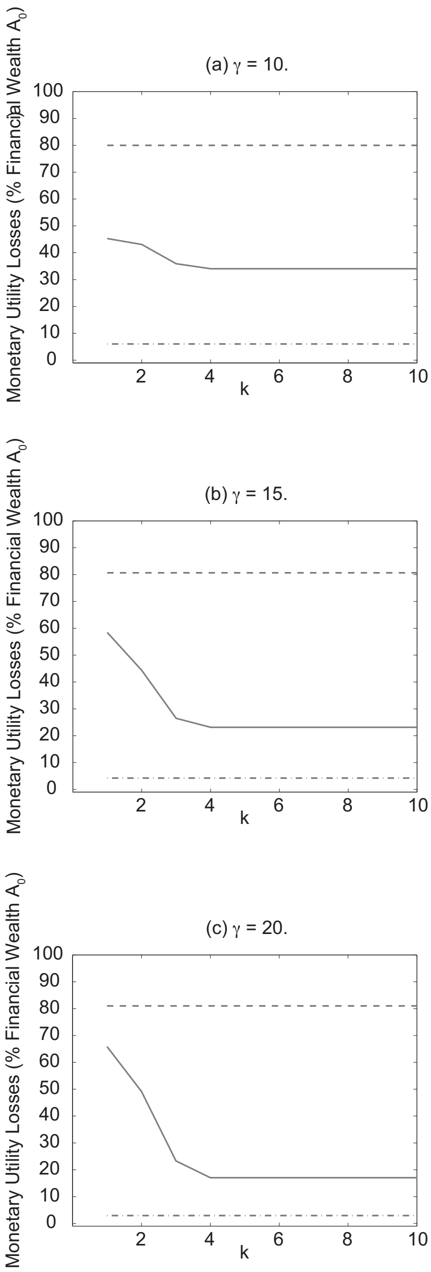

converges to the MUL for the optimal discretized strategy. The main difference with respect to figure 2 is that the MUL for the initial partition, which corresponds to a deterministic strategy based on a constant $\lambda^S = \bar{\lambda}$ value, is already low, since it is 7.6% for $\gamma = 20$ and 13.8% for $\gamma = 10$. This result has important practical implications, since it suggests that if for some reason the mutual fund industry was to restrict itself to purely deterministic strategies, the kind of deterministic strategies currently used could be substantially improved upon. On the other hand, restricting the set of strategies to purely deterministic strategies is not desirable as the marginal gain from introducing state-dependencies in the allocation strategy can become quite substantial for some reasonable parameter values. To illustrate this effect, we report in figure 6 the MUL for strategies based on partitions when the parameter λ^S exhibits more variability. Such a higher variability can be obtained by multiplying the standard deviation σ_λ by a factor of two, which results in multiplying the long-term mean volatility, $\sigma_\lambda / \sqrt{2\kappa}$, also by a factor of two. In fact the base case value of the long-term mean volatility $\sigma_\lambda / \sqrt{2\kappa} = 13.5\%$, was relatively low, and increasing it could very well provide a more accurate description of the actual underlying process. In this case, we then obtain that the MULs for the initial partition (corresponding to $k = 1$, which implies a deterministic strategy) range from 24.4% (for $\gamma = 20$) to 36.3% (for $\gamma = 10$), which indicates a substantial welfare loss from ignoring time-variation in the equity premium. The corresponding values in terms of return-to-risk ratio (see figure 7) also show substantial value to be added by switching from a deterministic strategy to a stochastic one that fully incorporates the impact of a time-varying opportunity set. When $k = 4$ (which corresponds to nine classes) on the other hand, the MUL obtained becomes lower than 10%. From figure 6, it can also be noted that the convergence of the MUL towards the discretized optimal strategy is slightly slower than in figure 4, since the quasi-equality between the two utility losses holds only as of $k = 5$. The same can be said of the return-to-risk ratio, as appears from figures 5 and 7. In unreported results we have checked that increasing the long-term variability of λ^S leads to higher MULs for the deterministic strategies obtained by letting $k = 1$. These findings confirm that the assumption of a constant Sharpe ratio involves a strong opportunity cost when the actual process exhibits significant variability.

c. Combined Impact of a Discrete Partition of the Set of Investors and Market Conditions

Having considered the separate effects of partitioning the set of investors and market conditions, we turn in figure 8 to the analysis of the joint impact of discretizing both dimensions. Roughly speaking, the final effect is the resultant of the individual effects of the two partitions. In the base case, the utility loss is a decreasing function of the index k and is indistinguishable from the MUL of the discretized optimal strategy for $k \geq 4$. Overall, our results suggest that even with a reasonably low total number of partitions, perfectly consistent with implementation constraints in a retail money management context, real world implementations of truly optimal strategies would represent a very significant improvement over the heuristic strategies currently used in the target date fund industry practice.

II.3. IMPACT OF AN IMPERFECT OBSERVATION OF THE EQUITY RISK PREMIUM

The previous strategies, including those based on partitions of the set of market conditions, rely on a perfect knowledge of the equity Sharpe ratio λ^S, which is a key input in the expression of the optimal strategy. In practice, however, the instantaneous expected return on equity indices is not directly observable. The literature on stock return predictability has proposed several proxies for the expected return, including notably an affine function of the dividend yield (see *e.g.* Campbell and Viceira (1999), Barberis (2000), Xia (2001) and Menzly *et al.* (2004)). However, the predictive power of the dividend-price ratio is rather weak: for example, regressing realized excess returns of stocks on the dividend yield, Campbell and Viceira (1999) find a R^2 of 2.8%; running a similar regression on an enlarged set of predictors, Campbell *et al.*(2003) find a R^2 of 8.6% only. This means that using such predictive regressions leaves a substantial amount of uncertainty over future equity returns. It is therefore hard to construct a reliable predictor of these returns. In this context, the natural question arises whether observing λ^S imperfectly has or not a strong impact on the superiority of optimal life-cycle investing strategies with respect to heuristic allocation schemes. As a first step towards addressing this problem, we conduct the following experiment: we assume that the agent observes a noisy version $\overline{\lambda}_t^S$, given by $\overline{\lambda}_t^S = \lambda_t^S + \varepsilon_t$, of the equity index Sharpe ratio. The error term ε_t is taken to be a white noise with mean zero (which means that the observation is unbiased) and variance denoted by σ_ε^2. A higher variance means a more noisy observation of λ^S. Then, life-cycle strategies are implemented with $\overline{\lambda}_t^S$ used as a (noisy) proxy for the true value of the Sharpe ratio.[13] Figures 10 and 11 show the MULs induced by strategies based on partitions of both the set of time-to-horizons and of the set of Sharpe ratios, assuming that a noisy version of the Sharpe ratio is observed. Two values of the standard deviation of the white noise ε are considered: a value equal respectively to twice and five time the long-term volatility of the Sharpe ratio, which corresponds to $\sigma_\varepsilon = 30\%$ and $\sigma_\varepsilon = 67.4\%$, respectively. As expected, we obtain that for any number of classes in the partition but $k = 1$, the MULs are higher than in the corresponding figure 8, in which the expected excess return was perfectly observed.[14] We also find that even after accounting for the presence of relatively high measurement error in the equity risk premium ($\sigma_\varepsilon = 67.4\%$), the MULs obtained are still substantially lower compared to those obtained for heuristic glide paths used in TDF practice.

A comparison of figures10 and 11 shows that the MUL is an increasing function of the variance of the noise, as was to be expected. More interestingly, the increase in the MUL that follows from introducing measurement error is found to be larger for high values of the index k. This finding can be explained by the fact that for low k values, the relative coarseness of the partition implies higher robustness since misclassification risk in this case is relatively limited. For $k = 2$ for example, there are only three admissible values, corresponding respectively to low, medium and high expected return levels, and in most cases the imperfect observation of the true value for the equity risk premium will not have any impact on which level (low, medium or high) should be used. For higher values of the index k, misclassification risk is a greater concern. In other words, while obtaining an accurate estimate for the equity risk premium is difficult, assessing whether the current value corresponds to a high, intermediate or low risk premium level is an easier task, and achieving such a more modest objective is sufficient for generating reasonably good proxies for optimal strategies. Overall, the superiority of optimal life-cycle investing strategies implemented with reasonably fine partitions with respect to heuristic glide paths currently used in the industry is found to be robust with respect to the introduction of reasonable levels of measurement errors in equity risk premium estimates.

IV. CONCLUSION

In this paper, we characterize in closed-form the optimal time – and state – dependent allocation strategy for a long-term investor preparing for retirement in the presence of stochastic interest and inflation rates and a mean-reverting equity risk premium. We find that the opportunity cost involved in purely deterministic life-cycle strategies such as those implemented by available *target date funds* is substantial for reasonable parameter values. Our results also suggest that a parsimonious partition of the set of investors and market conditions can be performed, which allows for a relatively accurate approximation of the true optimal strategy while being consistent with real-world implementation constraints in a retail money management context. Our results have important potential implications for the design of improved forms of target date funds. More financial innovation is critically needed at this stage to design better target date funds that could help investors coping with their long-term financial planning problems (see Bodie *et al.* (2009) for a related discussion). Overall, our results suggest that there is ample room for added value between one-(allocation)-size-fits-all (investors with same age) solutions and do-it-yourself approaches to life-cycle investment decisions.

Our work can be extended in a number of directions. On the one hand, we have made the simplifying assumption of a constant equity volatility. Incorporating time-variation in equity volatility, in addition to time-variation in equity returns, can be performed at the cost of a relatively modest increase in the model complexity (details are available from the authors upon request). Another important ingredient which is missing in our model, as well as in most of the literature on life-cycle investing, is how the presence of short-term performance constraints faced by investors would affect the optimal long-term allocation decisions. Incorporating explicit mechanisms for ensuring some protection against short-term downside equity risk would be a desirable ingredient for most long-term investors. Recent research provides useful insights in that direction, since it suggests that a general analytical representation of the relationship between optimal long-term strategies in the presence and in the absence of short-term constraints can be identified, which allows one to disentangle the impact of short-term constraints from the impact of return predictability on the optimal allocation decision (Martellini and Milhau, 2010b). ■

ACKNOWLEDGEMENTS

This research has benefited from the support of the UFG Life Cycle Investing and Target Date Funds research chair at EDHEC-Risk Institute. We would like to thank Noël Amenc, Xavier Lépine, Bacem Rezgui, as well as participants at the 2009 EDHEC Institutional Days seminar, for very useful comments. Any remaining error is ours.

1. One notable exception is Benzoni *et al.* (2007), who question the bond-like nature of human capital, and show that the presence of co-integration between labor income and stock dividends results in a high long-term correlation between stocks and human capital. This substantially reduces the optimal demand for stocks by young investors, and even results in short positions in the equity market.
2. A partial explanation for this could be related to the fact that TDF providers do not have a strong incentive to promote the use of low fee money market funds.
3. As such, our paper is closely related to Sangvinatsos and Wachter (2005).
4. Real estate assets can prove useful ingredients in long-term investment strategies both for their inflation-hedging benefits and for their diversification benefits with respect to stocks and bonds. The main messages in the paper, however, would hold true in the absence of a real estate asset.
5. Introducing income risk in this model is straightforward as long as this risk is spanned by traded assets. In the presence of unspanned income risk, however, the income stream can not be valued as the dividend flow of a traded asset/portfolio, and one has to introduce a utility-based investor-specific price (see Henderson (2005) or Munk (2000)).
6. This matrix is constant because we have assumed a constant maturity bond. If the nominal bond had a fixed maturity T_0, the volatility matrix would be time-dependent.
7. Overall, having a system of ordinary differential equations (ODEs) satisfied by the functions A_i greatly simplifies numerical simulations of the optimal strategy, which we shall use extensively in what follows. Closed-form expressions could possibly be obtained, but given their complexity they would arguably be of little help in understanding the impact of the various parameters (see *e.g.* munk2004daa for such expressions in a slightly different context, with two risky traded assets only).
8. Since A_4 is an increasing function of time-to-maturity, the term $A_4(T-t)\lambda_t^S$ tends to make the allocation to stocks an increasing function of horizon for investors with $\gamma > 1$ However, the final effect of an increase in investment horizon on the weight depends also on $A_3(T-t)$.
9. The choice of the maturity for the bond index has no impact, at least in the context of a continuous-time implementation, since a bond index with a given constant maturity is a perfect substitute for a bond index with any other given constant maturity.
10. The continuous-time parameters involved in our model could also be calibrated using a Kalman filter (see Harvey (1991)). This technique has been employed *e.g.* in Brennan and Xia (2002), Munk *et al.* (2004) and Sangvinatsos and Wachter (2005).
11. This imperfect partitioning of the set of investors should in principle also be performed as a function of risk-aversion in addition to time-horizon. On the other hand, because risk-aversion is not an observable parameter, measuring a welfare loss implied by a discrete partition is arguably not as straightforward as it is for an observable parameter such as time-horizon. In fact, investors will typically have to position themselves in various categories (say high, moderate and low risk-aversion levels), which then will be assumed to perfectly represent the investors' unobservable risk-aversion level, without necessarily incurring an explicit welfare cost related to an imperfect customization of the strategy according to this dimension. In the numerical analysis, we focus on the values $\gamma = 10, 15$, and 20. In practice, the choice of the risk-version parameter levels is typically calibrated so as to avoid undesirable extreme or short positions in the various asset classes.
12. Anari and Kolari (2002) house provide empirical evidence that house prices are a stable inflation hedge in the long-run.
13. Another problem is the presence of estimation risk on the parameters of the equity risk-premium process, as opposed to the realized value for the process.
14. For $k = 1$, the actual value of the Sharpe ratio is irrelevant since this stochastic quantity is treated as a constant equal to $\bar{\lambda}$.

References

- Anari A. and Kolari J. (2002). House prices and inflation. *Real Estate Economics* 30(1), 67-84.
- Barberis N. (2000). Investing for the long run when returns are predictable. *Journal of Finance* 55(1), 225-264.
- Basu A. and Brisbane A. (2009). Towards a dynamic asset allocation framework for target retirement funds: Getting rid of the dogma in lifecycle investing. Working Paper.
- Basu A. and Drew M. (2009). Portfolio size effect in retirement accounts: What does it imply for life-cycle asset allocation funds? *The Journal of Portfolio Management* 35(3), 61-72.
- Benzoni L., Collin-Dufresne P. and Goldstein R. (2007). Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *Journal of Finance* 62(5), 2123-2167.
- Bodie Z., Detemple J. and Rindisbacher M. (2009). Life cycle finance and the design of pension plans. Forthcoming in *Annual Review of Financial Economics*.
- Booth P. and Yakoubov Y. (2000). Investment policy for defined-contributionpension scheme members close to retirement: An analysis of the "lifestyle" concept. *North American Actuarial Journal* 4(2), 1-19.
- Brennan M., Schwartz E., and Lagnado R. (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control* 21(8-9), 1377-1403.
- Brennan M. and Xia Y. (2002). Dynamic asset allocation under inflation. *Journal of Finance* 57(3), 1201-1238.
- Cairns A., Blake D. and Dowd K. (2006). Stochastic lifestyling: Optimal dynamic asset allocation for defined contribution pension plans. *Journal of Economic Dynamics and Control* 30(5), 843-877.
- Campbell J., Chan Y. and Viceira L. (2003). A multivariate model of strategic asset allocation. *Journal of Financial Economics* 67(1), 41-80.
- Campbell J. and Viceira L. (1999). Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics* 114(2), 433-495.
- Cocco J., Gomes F. and Maenhout P. (2005). Consumption and portfolio over the life-cycle. *Review of Financial Studies* 18(2), 491-533.
- Dai Q. and Singleton K. (2000). Specification analysis of affne term structure models. *Journal of Finance* 55(5), 1943-1978.
- DeJong F., Driessen J. and Van Hemert O. (2007). Hedging house price risk: Portfolio choice with housing futures.
- Detemple J., Garcia R. and Rindisbacher M. (2003). A Monte Carlo method for optimal portfolios. *Journal of Finance* 58(1), 401-446.
- Harvey A. (1991). *Forecasting, structural time series models and the Kalman filter*. Cambridge Univ Press.
- He H. and Pearson N.D. (1991). Consumption and portfolio policies with incomplete markets and short-sale constraints: The infinite dimensional case. *Journal of Economic Theory* 54(2), 259-304.
- Henderson V. (2005). Explicit solutions to an optimal portfolio choice problem with stochastic income. *Journal of Economic Dynamics and Control* 29(7), 1237-1266.
- Kim T. and Omberg E. (1996). Dynamic nonmyopic portfolio behavior. *Review of Financial Studies* 9(1), 141-161.
- Lewis N., Okunev J. and White D. (2007). Using a Value-at-Risk approach to enhance tactical asset allocation. *The Journal of Investing* 16(4), 100-107.
- Lioui A. and Poncet P. (2001). On optimal portfolio choice under stochastic interest rates. *Journal of Economic Dynamics and Control* 25(11), 1841-1865.
- Liu J. (2007). Portfolio selection in stochastic environments. *Review of Financial Studies* 20(1), 1-39.
- Martellini L. and Milhau V. (2010a). From deterministic to stochastic life-cycle investing - implications for the design of new forms of target-date funds. EDHEC-Risk Institute Publication.
- Martellini L. and Milhau V. (2010b). Long-term investing with short-term constraints. Working Paper, EDHEC-Risk Institute.
- Menzly L., Santos T. and Veronesi P. (2004).Under standing predictability. *Journal of Political Economy* 112(1).
- Merton R. (1971). Optimal portfolio and consumption rules in a continuous-time model. *Journal of Economic Theory* 3(4), 373-413.
- Mitchell O., Mottola G., Utkus S. and Yamaguchi T. (2006). The Inattentive Participant: Portfolio Trading Behavior in 401(k) Plans. *Pension Research Council* 2006-2.
- Munk C. (2000). Optimal consumption/investment policies with undiversifiable income risk and liquidity constraints. *Journal of Economic Dynamics and Control* 24(9), 1315-1343.
- Munk C. and Sørensen C. (2004). Optimal consumption and investment strategies with stochastic interest rates. *Journal of Banking and Finance* 28(8), 1987-2013.
- Munk C. and Sørensen C. (2010). Dynamic asset allocation with stochastic income and interest rates. *Journal of Financial Economics*. Forthcoming.
- Munk C., Sørensen C. and Vinther T. (2004). Dynamic asset allocation under mean-reverting returns, stochastic interest rates and inflation uncertainty: Are popular recommendations consistent with rational behavior? *International Review of Economics and Finance* 13, 141-166.
- Sangvinatsos A. and Wachter J. (2005). Does the failure of the expectations hypothesis matter for long-term investors? *Journal of Finance* 60(1), 179-230.
- Shiller R. (2005). The life-cycle personal accounts proposal for social security: An evaluation. *NBER Working Paper* (11300).
- Sørensen C. (1999). Dynamic asset allocation and fixed income management. *Journal of Financial and Quantitative Analysis* 34(4), 513-531.
- Strategic Insight (2009). Guaranteed retirement income beyond annuities. Available at http://www.annuityinsight.com/annuityinsight/published/2009-SALB/main.aspx.
- Towers Watson (2010). 2010 global pension asset study. Available at http://www.towerswatson.com/assets/pdf/966/GPAS2010.pdf.
- Vasicek O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics* 5(2), 177-188.
- Viceira L. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *Journal of Finance* 56(2), 433-470.
- Viceira L. and Field S. (2007). Life-cycle funds.
- Wachter J. (2002). Optimal consumption and portfolio allocation under mean-reverting returns: An exact solution for complete markets. *Journal of Financial and Quantitative Analysis* 37, 63-91.
- Wise D. A. (Ed.) (2009). *Developments in the Economics of Aging*, pp. 15-56. University Of Chicago Press.
- Xia Y. (2001). Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation. *Journal of Finance*, 205-246.

CFA Institute—EDHEC-Risk Institute

Advances in Asset Allocation Seminar

State-of-the-art risk management and asset allocation processes for product innovation

London, 30 November—2 December 2010

New York, 8—10 March, 2011

Singapore, 29—31 March 2011

Bridging the Gap between Portfolio Theory and Portfolio Construction

Optimising Portfolio Construction in a Non-Markowitz World

Exploring New Forms of Indices and Benchmarks

Incorporating Alternative Investments into ALM for Diversification, Substitution, and Inflation Hedging

Integrating Risk Management Constraints in Asset Allocation

Designing Long-Only Absolute Return Funds and New LDI Solutions

"*A must for all portfolio managers, and client advisors working on either asset allocation or LDI implementation. In this time of resource constraints, it is critical to understand where to focus your research efforts and this seminar is instrumental in understanding how to take the advances in available technology and improve upon portfolio construction and risk management.*

Anwiti Bahuguna, Director, Bank of America and
Senior Portfolio Manager, Columbia Management, USA
Past participant

Excellent seminar bringing together the latest in risk management with better asset management solutions.

Alwin Oerlemans, CFA
Director Client Solutions, APG Institutional Clients, The Netherlands
Past participant

An excellent support for private wealth management.

Diego Caballero Yañez, CFA
Investments Specialist, BSI, Switzerland
Past participant

Probably the best learning experience I have had on the topic.

Jose Vidrago, Head of Research & Operations, IGFCSS, Portugal
Past participant

Lionel Martellini, PhD
Scientific Director,
EDHEC-Risk Institute

Information and registration:
Mélanie Ruiz: +33 (0) 493 187 819

EXECeducation@edhec-risk.com

http://store.edhec-risk.com

CFA INSTITUTE

www.cfainstitute.org

EDHEC-RISK Institute

www.edhec-risk.com