

Evaluation of target date funds

Problem Presenter: Scott Warlow (Manulife Financial)

Academic Participants: Mirela Cara (York University), Myriam Cisneros-Molina (Mexican Petroleum Institute), David Cottrel (McGill University), Huaxiong Huang (York University), Jiawei Li (York University), Ashley Pitcher (University of Oxford), Yun Qiao (York University), Thomas S. Salisbury (York University)

With help from: Joel Phillips (McGill University), and Pavan Aroda (York University)

Report prepared by: Myriam Cisneros-Molina, Huaxiong Huang and Thomas S. Salisbury¹

Abstract. Target date funds are an emerging class of investment products, designed for retirement savings. The project considered methodologies for ranking such funds.

1 Introduction

At the 2008 Fields-MITACS Industrial Problems Workshop, Scott Warlow from Manulife’s strategic asset allocation group presented a problem to the workshop participants. In the past ten years, a special extension of the fund-of-fund mutual funds, known as target-date-funds (TDFs) has become available for pension funds in the US and Canada. These funds bundle several mutual funds together, and target a specific exposure to stocks and bonds. Unlike most fund-of-fund products, TDFs continuously adjust the equity weight from a high level, e.g., 80-90% 40 year before the target date, to a lower level, e.g., 40-50% at the target date. Some continue to reduce the equity weight into retirement. The Pension Protection Act of 2006 in the US redefined qualified default investment alternatives for defined contribution (DC) pension plans to include TDFs, which has led to a great deal of interest in these funds by DC plan sponsors [7].

There have been several attempts by practitioners and academics to introduce a methodology to rank TDFs. Some of them use approaches adopted from evaluating conventional mutual funds. There are several issues associated with these approaches. The main issue is that they usually do not take into account the special characteristics of the TDFs. For

¹hhuang@yorku.ca, and salt@yorku.ca

example, methods for evaluating mutual funds are typically one-period, on a yearly basis. TDFs, on the other hand, are multi-period in nature, over a long time horizon. In addition, TDFs usually have two phases: pre-retirement (contribution) and post-retirement (withdrawal). Since they cover the full life cycle of the individual, all forms of wealth are important, and human capital plays a pivotal role. These issues make the evaluation and ranking of TDFs more difficult than the conventional mutual funds.

Historical data has been used in combination of a block-bootstrap method to evaluating TDFs by Labovitz [5], focusing on the distribution of wealth at retirement. Clark [2] and Turowski [6] used the Bradley-Terry model for paired comparisons, or paired comparison digraphs to rank TDFs. Industry groups, such as Dow Jones, Morningstar, or Turnstone Advisory Group, have tried to produce benchmarks for glide paths, peer categories, or a score function. The main shortcomings of these approaches are that the categories are either very broad, or the focus is very narrow (for the accumulation phase anyway). In general, the unique characteristics of TDFs are not considered by these approaches.

In a recent paper, Bodie and Treussard [1] proposed a method in the spirit of indifference pricing to evaluate TDFs using a continuous time optimal control framework. Given an investor’s level of risk aversion and labor income, the optimal strategy generates a combination of the TDF and a risk-free bond. The return of the TDF is compared with the return of the optimal strategy and serves as an indicator for evaluating the TDF. While their assumption on the labor income is quite restrictive, as is the focus on the accumulation phase, the approach used in [1] is attractive since it offers an unbiased way to rank TDFs.

During the workshop, the participants explored the ideas of using indifference pricing and discussed the possibility of extending the Bodie and Treussard approach. In particular, the assumption on labor income can be relaxed, using ideas from [3, 4]. Furthermore, an extension of the approach in [1] to both accumulation and withdrawal phases can also be carried out using the method in [3, 4], where insurance is also included in developing the optimal strategy. This can be done by incorporating labour income into the classical optimal asset allocation problem as in [3, 4]. Having obtained the “ideal” glide path for an individual, we can then in principle compare and rank TDF’s by comparing their expected utility to that of an optimal portfolio. In this report, we will outline such an approach and provide preliminary numerical results.

Scott Warlow formulated the problem, gave the initial presentation at the beginning of the workshop, and gave feedback at its end. Tom Salisbury and Huaxiong Huang advised the project students during the project week, and prepared the final version of the report. Mirela Cara, Myriam Cisneros-Molina, David Cottrel, Jiawei Li, Ashley Pitcher, and Yun Qiao all worked intensively on the project during the workshop, with periodic input from Joel Phillips. Mirela Cara, Myriam Cisneros-Molina and Yun Qiao prepared a presentation for workshop participants. David Cottrel provided numerical results and plots for the presentation and report. Myriam Cisneros-Molina drafted the initial report, with input from Yun Qiao on the mathematical derivations. Pavan Aroda conducted a literature search, after the workshop ended.

2 Optimization and stochastic control

A balanced fund is a mutual fund with a specified asset allocation. In other words, balanced funds fix the percentage they allocated to stocks and to bonds. The classical Merton asset allocation problem suggests that such a fixed asset allocation should be optimal for many individuals. But this classic problem neglects many types of wealth, other than

stocks and bonds. In particular, it neglects labour income, which for many individuals is predictable and smooth, up to retirement. In other words, labour income can be viewed as equivalent to the coupon stream from a bond, with maturity the individual’s retirement date. Incorporating that into the Merton problem should therefore perturb the optimal asset allocation, so that the allocation of pension savings to stocks is high initially (when the labour income “bond” has high value), and then drops (as the labour income “bond” decreases in value). Target date funds do precisely this.

Target date funds are pension savings products that include allocations to stocks and to bonds. Unlike a balanced fund, that stock allocation changes over time. Individuals saving for retirement, who are unable or uncomfortable managing their asset allocation actively (or who would find it expensive to do so) can therefore use a target date fund to achieve outcomes that are closer to optimal than a fixed-allocation investment strategy.

Each fund has its own way of adjusting the allocation over time, called a *glide path* – see Figure 1

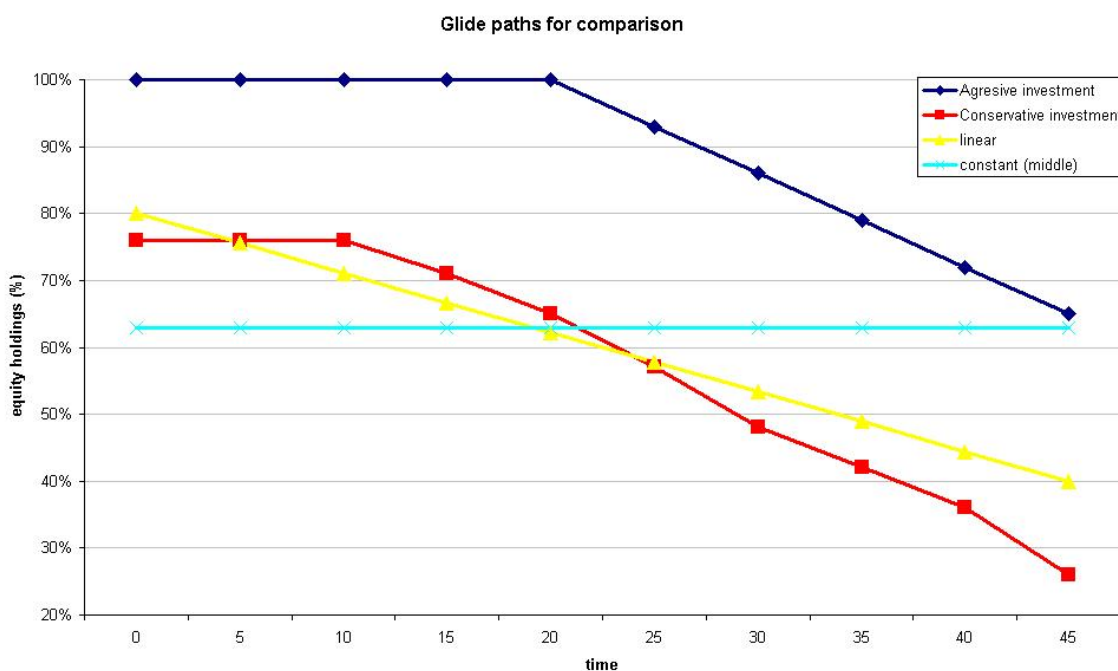


Figure 1 Simple glide paths

Our basic approach to evaluating a TDF is to compute an individual’s optimal asset allocation over time, and compare it to the glide path of the TDF. One simple basis for comparison would simply be to see what glide path best fits the computed allocation best. The variety of glide paths available is in fact quite rich, so in many cases there will be a good fit available – see Figure 2.

A more quantitative approach, suitable for ranking TDFs, according to their suitability for a particular individual, would be to compare the expected utility achieved by the glide path to the expected utility achieved by the optimal asset allocation curve. We illustrate this approach below. If the expected utility achieved by a particular TDF is close to optimal,

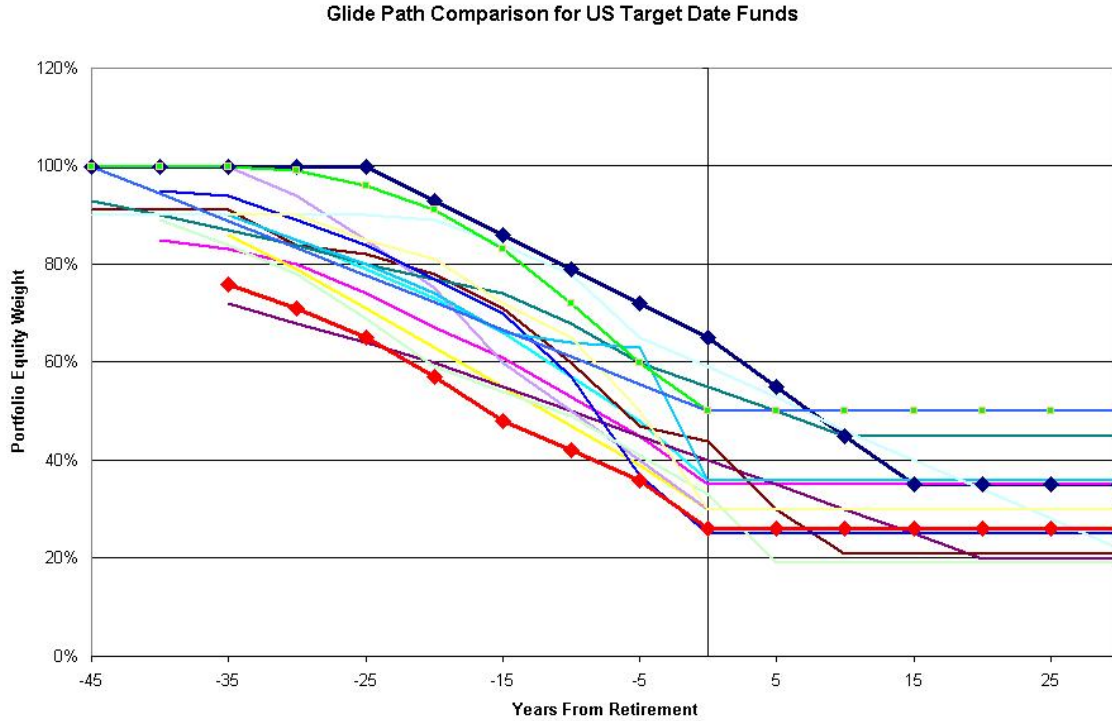


Figure 2 Sampling the universe of glide paths

then convenience of making a single purchase, and leaving all investment and allocation decisions to the vendor, may make this a suitable purchase for many consumers. Especially if they do not have the expertise or desire to actively supervise their portfolio, and to adjust the asset allocation over time. Depending on the fees involved, this may well be a cost-effective way of implementing a near-optimal retirement savings plan.

The optimal allocation curve will depend on both the consumer's risk preferences and the statistical properties of their income stream. In reality, neither risk preferences nor income are known with any precision, making the problem of comparing different target date funds more complicated than the version of this problem that we consider. An even broader problem would be that faced by a company considering choosing a target date fund as a vehicle for that company's pension plan. In that case, one could adopt a similar approach, by aggregating risk preferences and income paths across a group of employees. We do not consider that problem, but it is a natural extension of our approach.

Turning to the optimal asset allocation problem, we assume there exists a risk-free asset (a government bond) which grows exponentially over time at a fixed rate:

$$dR_t = rR_t dt.$$

We assume that the investment in equity is captured by trading in a single diversified equity portfolio S_t (e.g. a stock index such as the S&P 500) whose value (under the physical/real

world probability measure P) follows a geometric Brownian motion:

$$dS_t = S_t [\mu dt + \sigma dB_t].$$

Here $\mu > 0$ is the drift, and $\sigma > 0$ is the volatility, both assumed to be constant. B_t is a Brownian motion under the measure P .

Remark 2.1 A possible extension of our work would be to allow stochastic interest rates and stochastic equity volatility. Since pension savings involve investments over a long time horizon, adding such effects could well be important when doing realistic comparisons. Likewise, inflation is a significant risk over such time periods, so one would like to model inflation and incorporate that into the problem. We leave such questions for another day.

In order to capture the uncertainty of future labour income, the individual consumer's rate of wage earnings evolves randomly over time according to another geometric Brownian motion:

$$de_t = e_t [\mu_e dt + \sigma_e dB_t],$$

subject to initial earnings rate e_0 . In this first version of the problem, we are taking S_t and e_t to be perfectly correlated.

If π_t represents the fraction of wealth invested in the risky asset S_t , and labour earnings flow into the retirement account at a rate $e_t dt$, we have a current wealth process which evolves according as follows:

$$dX_t = e_t dt + X_t r dt + X_t \pi_t (\mu - r) dt + X_t \pi_t \sigma dB_t,$$

together with an initial condition $X_0 = x$. The terminology (wealth, labour income) suggests a broad interpretation of this analysis. For purposes of comparing TDF's, we take a narrower view, and regard X_t as the value of a retirement savings plan, with e_t representing the rate contributions to that plan.

We represent the consumer's preferences by a constant relative risk aversion (CRRA) utility function for terminal wealth

$$U(X_T) = \frac{X_T^{1-\gamma}}{1-\gamma}.$$

Here γ is a relative risk aversion parameter, and T is interpreted as the retirement date. Thus, the problem consists of finding the optimal policy π_t^* that maximizes utility from terminal wealth. In other words, our value function is

$$v(t, x, e) = \max_{\pi} \mathbb{E} \left[\frac{X_T^{1-\gamma}}{1-\gamma} \mid X_t = x, e_t = e \right].$$

We can exploit the scaling symmetry of our problem, to reduce its dimension. In particular, if we replace X_t by cX_t and e_t by ce_t , the dynamics remain the same. Thus

$$v(t, cx, ce) = c^{1-\gamma} v(t, x, e).$$

Letting $u(t, y) = v(t, y, 1)$ we get that $v(t, x, e) = e^{1-\gamma} u(t, x/e)$.

With that reduction, we turn to finding the Hamilton-Jacobi-Bellman (HJB) equations for $u(t, y)$. To do that, define $Y_t = X_t/e_t$. The dynamics of Y_t are that

$$\begin{aligned} dY_t &= \frac{dX_t}{e_t} - \frac{X_t de_t}{e_t^2} + \frac{X_t d\langle e \rangle_t}{e_t^3} - \frac{d\langle X, e \rangle_t}{e_t^2} \\ &= [1 + Y_t(r + (\mu - r)\pi_t)] dt + Y_t\pi_t\sigma dB_t \\ &\quad - Y_t[\mu_e dt + \sigma_e dB_t] + \sigma_e^2 Y_t dt - \sigma\sigma_e\pi_t Y_t dt \\ &= [1 + Y_t(r + \sigma_e^2 - \mu_e + (\mu - r - \sigma\sigma_e)\pi_t)] dt + (\pi_t\sigma - \sigma_e) Y_t dB_t. \end{aligned}$$

If there is indeed an optimal control π_t^* , then using it as control, $v(t, X_t^*, e_t)$ will be a conditional expectation, and hence a martingale. And if we use an arbitrary control π_t , there will be slippage over time in the expected utility. In other words, $v(t, X_t, e_t)$ will decrease (on average) over time: $\mathbb{E}[v(t, X_t, e_t) | \mathcal{F}_s] \leq v(s, X_s, e_s)$ for $s < t$. That is, $v(t, X_t, e_t)$ is a supermartingale. We can therefore obtain the HJB equation by computing $de_t^{1-\gamma}u(t, Y_t)$ using Ito's lemma. It will have the form $A_t dt + H_t dB_t$ for some A and H . The martingale property means that $A = 0$ when the optimal control π_t^* is used. The supermartingale property means that $A \leq 0$ when any other control is used. In other words, $\sup_\pi A = 0$.

Carrying this procedure out, we obtain the following equation:

$$\begin{aligned} 0 = u_t + \left[(1 - \gamma)\mu_e - \frac{1}{2}\gamma(1 - \gamma)\sigma_e^2 \right] u + [1 + y(r - \mu_e + \gamma\sigma_e^2)] u_y + \\ + \max_{\pi_t} \left[(\mu - r - \gamma\sigma\sigma_e)\pi_t y u_y + \frac{1}{2}(\pi_t\sigma - \sigma_e)^2 y^2 u_{yy} \right]. \end{aligned}$$

Provided $u_{yy} < 0$, we obtain the π that achieves the maximum, by differentiating a quadratic expression in π . That is,

$$\pi^*(t, y) = -\frac{(\mu - r - \gamma\sigma\sigma_e)}{\sigma^2} \frac{u_y}{y u_{yy}} + \frac{\sigma_e}{\sigma},$$

and the above max becomes

$$-\frac{(\mu - r - \gamma\sigma\sigma_e)^2 u_y^2}{2\sigma^2 u_{yy}} + \frac{\sigma_e}{\sigma} (\mu - r - \gamma\sigma\sigma_e) y u_y.$$

Substituting, the value function $u(t, y)$ should satisfy the following nonlinear PDE

$$\begin{aligned} 0 = u_t + \left[(1 - \gamma)\mu_e - \frac{1}{2}\gamma(1 - \gamma)\sigma_e^2 \right] u - \frac{1}{2} \frac{(\mu - r - \gamma\sigma\sigma_e)^2}{\sigma^2} \frac{u_y^2}{u_{yy}} + \\ + \left[1 + y \left(r - \mu_e + \frac{\sigma_e}{\sigma} (\mu - r) \right) \right] u_y. \end{aligned}$$

We also have the terminal condition

$$u(T, y) = \frac{y^{1-\gamma}}{1-\gamma}.$$

Taking inspiration from solutions to portfolio optimization problems [4] under HARA utility, we seek a solution of the form

$$u(t, y) = \frac{(y + \alpha(t))^{1-\gamma}}{1-\gamma} \beta(t).$$

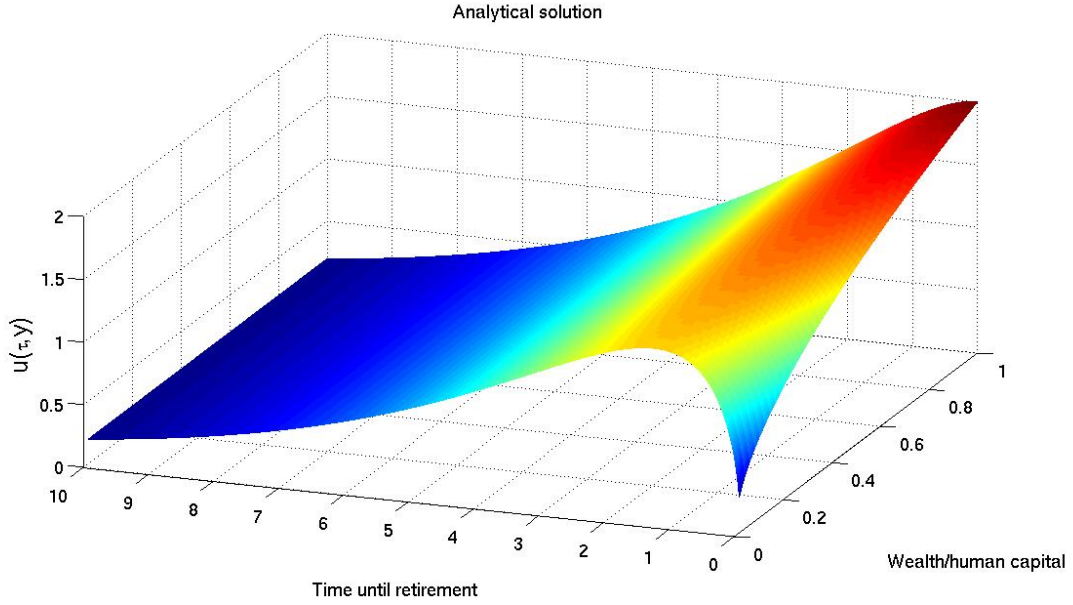


Figure 3 Optimal $u(t, y)$, with perfect correlations

Rewrite the PDE as

$$0 = u_t + a_0 u - a_2 \frac{u_y^2}{u_{yy}} + [1 + a_1 y] u_y,$$

where

$$\begin{aligned} a_0 &= (1 - \gamma)\mu_e - \frac{1}{2}\gamma(1 - \gamma)\sigma_e^2, \\ a_1 &= r - \mu_e + \frac{\sigma_e}{\sigma}(\mu - r), \quad \text{and} \\ a_2 &= \frac{1}{2} \frac{(\mu - r - \gamma\sigma\sigma_e)^2}{\sigma^2}. \end{aligned}$$

Substituting, we have the equation

$$0 = \frac{\beta'}{\beta} + a_0 + a_1(1 - \gamma) + a_2 \left(\frac{1 - \gamma}{\gamma} \right) + \frac{1 - \gamma}{y + \alpha} [\alpha' + 1 - a_1 \alpha].$$

It follows that we obtain a solution if we take

$$\begin{aligned} \alpha(t) &= \frac{1}{a_1} (1 - \exp\{-a_1(T - t)\}), \quad \text{and} \\ \beta(t) &= \exp \left\{ \left(a_0 + a_1(1 - \gamma) + a_2 \left(\frac{1 - \gamma}{\gamma} \right) \right) (T - t) \right\}. \end{aligned}$$

See Figure 3 for a graphical representation of this solution.




Figure 4 Optimal $t \mapsto \pi^*(t, y)$, for various $y = x/e$ (perfect correlations)

Then, the optimal strategy will be expressed as

$$\begin{aligned} \pi^*(t, y) &= \frac{(\mu - r - \gamma\sigma\sigma_e)}{\gamma\sigma^2} \left(1 + \frac{\alpha(t)}{y} \right) + \frac{\sigma_e}{\sigma} \\ &= \frac{(\mu - r - \gamma\sigma\sigma_e)}{\gamma\sigma^2} \left(1 + \frac{1}{a_1 y} [1 - \exp\{-a_1(T-t)\}] \right) + \frac{\sigma_e}{\sigma}. \end{aligned}$$

At $t = T$ all allocations converge to the value

$$\pi^*(T, y) = \frac{\mu - r}{\gamma\sigma^2},$$

which is independent of the wealth-to-salary value y . See Figure 4, which uses parameters $\mu = 0.08$; $\mu_e = 0.035$; $r = 0.03$; $\gamma = 3$; $\sigma = 0.2$; $\sigma_e = 0.05$; $T = 45$; Note that the bigger y is, the more aggressive the allocation to equity appears to be.

We can carry out a similar analysis if the processes S_t and e_t are not perfectly correlated. In that case our model becomes

$$dS_t = S_t [\mu dt + \sigma dB_t^1], \quad de_t = e_t [\mu_e dt + \sigma_e dB_t^2],$$

where B_t^1 and B_t^2 are Brownian motions with $d\langle B^1, B^2 \rangle_t = \rho dt$, so ρ is the correlation coefficient. We get

$$dY_t = \left[1 + Y_t (r + \sigma_e^2 - \mu_e + (\mu - r - \rho\sigma\sigma_e)\pi_t) \right] dt + Y_t (\pi_t\sigma dB_t^1 - \sigma_e dB_t^2).$$

Then $d\langle Y \rangle_t = Y_t^2(\pi_t^2\sigma^2 + \sigma_e^2 - 2\rho\pi_t\sigma\sigma_e) dt$ and $d\langle Y, e \rangle_t = Y_t\sigma_e e_t(\rho\pi_t\sigma - \sigma_e) dt$, so the HJB equation becomes

$$0 = u_t + \left[(1 - \gamma)\mu_e - \frac{1}{2}\gamma(1 - \gamma)\sigma_e^2 \right] u + [1 + y(r - \mu_e + \gamma\sigma_e^2)] u_y + \\ + \max_{\pi_t} \left[(\mu - r - \rho\gamma\sigma\sigma_e) \pi_y u_y + \frac{1}{2} (\pi^2\sigma^2 + \sigma_e^2 - 2\rho\pi\sigma\sigma_e) y^2 u_{yy} \right].$$

Optimizing over π gives

$$\pi^*(t, y) = -\frac{(\mu - r - \rho\gamma\sigma\sigma_e)}{\sigma^2} \frac{u_y}{y u_{yy}} + \frac{\sigma_e}{\sigma} \rho.$$

Substituting back, the value function satisfies the following non-linear PDE:

$$0 = u_t + \left[(1 - \gamma)\mu_e - \frac{1}{2}\gamma(1 - \gamma)\sigma_e^2 \right] u - \frac{1}{2} \frac{(\mu - r - \rho\gamma\sigma\sigma_e)^2}{\sigma^2} \frac{u_y^2}{u_{yy}} + \\ + \frac{1}{2} \sigma_e^2 (1 - \rho^2) y^2 u_{yy} + \left[1 + y \left(r - \mu_e + \gamma\sigma_e^2 + (\mu - r - \rho\gamma\sigma\sigma_e) \frac{\sigma_e}{\sigma} \rho \right) \right] u_y,$$

with the same terminal condition $u(T, y) = y^{1-\gamma}/(1-\gamma)$.

Remark 2.2 The analysis presented here is similar to that in [3, 4]. In fact, students working on the project needed to spend significant time mastering the techniques of stochastic optimal control. There are differences in the details of this analysis, compared to those of [3, 4]. First of all, since our goal to compare various TDFs, we ignore extraneous factors, such as insurance and consumption. Instead we focus on a pension or retirement-savings plan, and think of the labour income included in the analysis as modelling the employees' contribution to the pension funds. Secondly, and for the same reason, the analysis presented in this report is for the accumulation phase only. Finally, we have not carried the analysis of the correlated problem further, or extended the utility function to the HARA class, since in general one then has to use numerical methods to solve the HJB equations. We leave this for further work. See [3, 4] for further details.

3 Monte Carlo analysis

Having obtained the optimal asset allocation over time, we now need to compare the expected utility using a TDF glide path to the optimal expected utility. We carried this out in the perfectly correlated case. To do so, we simulated the paths of both e_t and X_t – see Figure 5 for these and for paths of Y_t computed from e_t and X_t . One can also compute the path of the optimal allocation $\pi^*(t, Y_t)$. See Figure 6, in which the paths are traced out along the surface of optimal $\pi^*(t, y)$.

We then computed the realized utility of terminal wealth, for each simulation, using one of two glide paths (aggressive and conservative). We compute the mean over the simulations, and then compute a ratio, with that as the denominator and the optimal expected utility as the numerator. The results are shown in Figure 7.

There are two modifications to this procedure that would be desirable, before using this as a practical method of comparison. The first derives from the observation that the ratios




Figure 5 Simulations of e_t , X_t , and Y_t , plotted against $t - T$

found appear to be close to 1. To understand whether the observed differences are actually meaningful, it would be better to express the comparison in dollar terms, rather than in terms of dimensionless ratios of utilities. In other words, to compute the dollar premium that one would need to add to one's TDF portfolio, in order that its expected utility should match that of the optimal portfolio. That would put the comparison in terms of the same units (dollars) as the portfolios themselves. Consumers would then be positioned to decide whether the differences between TDFs are meaningful, in monetary terms.

A second modification would be to attempt to derive analytic expressions for the TDF expected utility. Or at least, to derive a PDE for this utility and then solve it numerically. We leave both of these extensions for future work.

4 Conclusion

In this report, we have outlined an approach based on indifference pricing for evaluating target-date-funds (TDFs). Our approach is based on previous work [1, 4]. This approach uses a continuous-time optimal control framework. When an investor's risk aversion and labor income process are known, an optimal asset allocation strategy between the underlying stock fund and a risk-free bond can be found. The relative ranking of the TDF can be obtained by computing the ratio of the expected utility under the TDF with that of the




Figure 6 Simulation of $\pi^*(t, Y_i)$ on the surface $w = \pi^*(t - T, y)$

optimal strategy. Comparing to some of the other approaches used in the industry, this approach is attractive as it provides an unbiased ranking which is tailored for individual clients.

One possibility for future work would be to expand the model to incorporate the retirement phase, the effects of the correlation between the labor income and the risky asset, stochastic interest rates, volatility, and inflation (all important effects during the accumulation phase for retirement savings). Another would be to evaluate the impact on utility of a fund manager's active management of the glide path (rather than adhering to a rigid formula). Finally, we note that in practice there are more than one risky asset in a TDF. An extension to the multi-asset case should be straightforward. An alternative approach, is to replace the risky asset in our analysis by the TDF, and solve the optimal asset allocation problem between the TDF and a risk-free bond, as in [1], for both the accumulation and retirement phases.

References

- [1] Bodie, Z. and Treussard, J. *Making Investment Choices as Simple as Possible, but Not Simpler*. Financial Analysts Journal **63** (2007), 42–47.
- [2] Clark, A. *Ranking Target-Date Funds Using Past Performance*. Lipper Research Series Fund Industry Insight Report (2007).

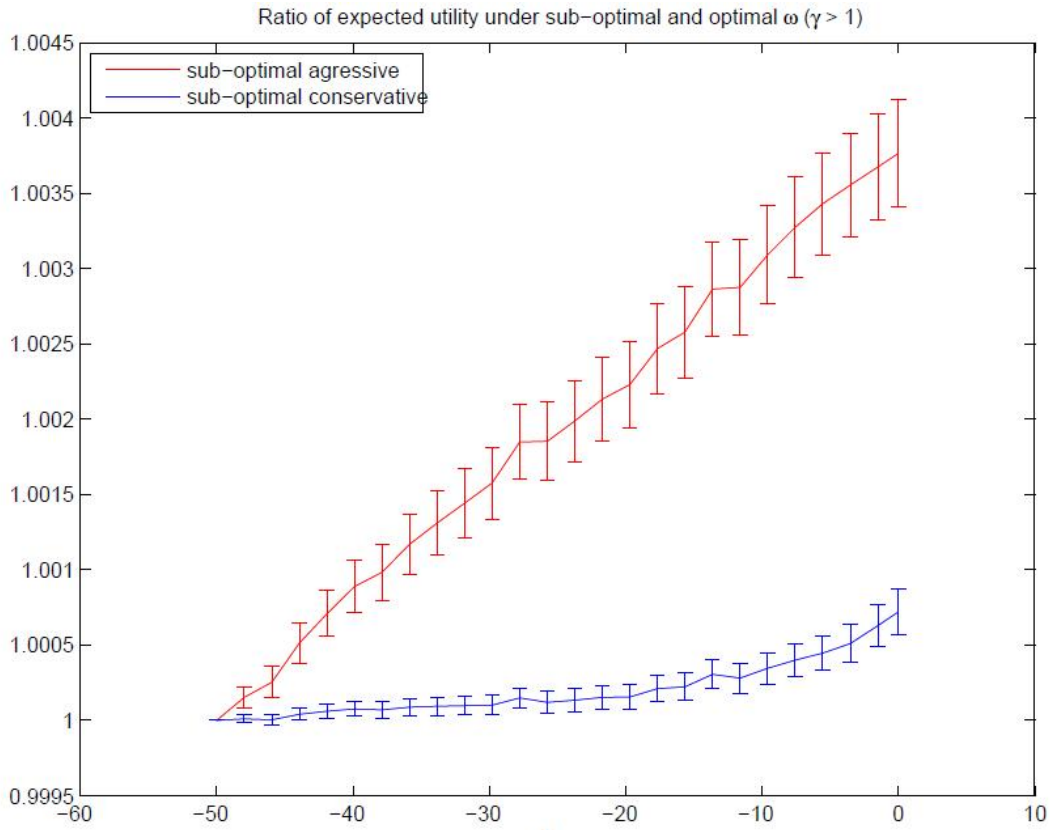


Figure 7 Expected optimal utility / Expected TDF utility

- [3] Huang, H., Milevsky, M.A., and Wang, J. *Portfolio Selection and Life Insurance*. Journal of Risk and Insurance **75** (2008), 848–872.
- [4] Huang, H. and Milevsky M.A. *Portfolio Choice and Mortality-Contingent Claims: The General HARA Case*. Journal of Banking and Finance **32** (2008), 2444–2452.
- [5] Labovitz, M.L. *Targeting Target Maturity Funds*. Lipper Research Series Fund Industry Insight Report (2006).
- [6] Turowski, H. *Using Paired Comparison Digraphs to Rank Target-Date Funds*. Lipper Research Series FundIndustry Insight Report (2007).
- [7] Warlow, S. *The Evaluation of Target-Date Mutual Funds*. (2008)
www.fields.utoronto.ca/audio/08-09/FMIPW/warlow/