Tactical Target Date Funds

Francisco Gomes* Alexander Michaelides† Yuxin Zhang‡

February 2018

*Department of Finance, London Business School, London NW1 4SA, UK. E-mail: fgomes@london.edu.
†Department of Finance, Imperial College London, South Kensington Campus, London SW7 2AZ, UK. E-mail: a.michaelides@imperial.ac.uk.
‡Hanqing Advanced Institute of Economics and Finance, RenMin University, Beijing, China. E-mail: yx.zhang@ruc.edu.cn.

Tactical Target Date Funds

Abstract

Tactical target date funds (TTDFs) modify conventional life style TDFs to take advantage of predictability in excess returns driven by the variance risk premium. In the context of a standard life-cycle model of consumption and savings decisions we show that investing in TTDFs generates economically large welfare gains. This result holds despite the portfolio rule followed by these funds is designed to be extremely simplified to be easily implementable and communicated to investors, and therefore significantly misspecified relative to the optimal policy functions. Crucially, the welfare gains remain large even after we introduce turnover restrictions that limit the annual turnover of the TTDFs to be comparable to that of the average mutual fund, and after we take into account potential increases in transaction costs. These gains are particularly higher for more risk tolerant investors.

JEL Classification: G11, D14, D15

Key Words: Target date funds, life cycle portfolio choice, variance risk premium, strategic asset allocation, tactical asset allocation, market timing.

1 Introduction

The conventional financial advice is that households should invest a larger proportion of their financial wealth in the stock market when young and gradually reduce the exposure to the stock market as they grow older. This is advice is given by several financial planning consultants (for instance, Vanguard) who recommend target-date funds (TDFs) that reduce exposure to the stock market as retirement approaches. The long term investment horizon in these funds, and the slow decumulation of risky assets from the portfolio as retirement approaches, can be thought of as strategic asset allocation (see Campbell and Viceira, 2002), where a long term objective (financing retirment) is optimally satisfied through the TDF. This investment approach arises naturally in the academic literature in the presence of undiversifiable labor income risk (for example, Cocco, Gomes, and Maenhout (2005), Gomes and Michaelides (2005), Polkovnichenko (2007), and Dahlquist, Setty and Vestman (forthcoming)).[1] Moreover, the most recent empirical evidence shows that, even outside of these pension funds, households follow this life-cycle investment pattern (Fagereng, Gottlieb and Guiso (2017)).

In this paper we investigate how exploiting time variation in expected returns can enhance the strategic asset allocation perspective of a life cycle investor saving for retirement, through tactical asset allocation movements over a quarterly frequency.[2,3] More precisely we consider a recently proposed predictability factor, the variance risk premium (hereafter VRP) proposed by Bollerslev, Tauchen and Zhou (2009) and Bollerslev, Marrone, Xu, and Zhou (2014)). Crucially, we explore how the welfare gains from the optimal policies can be

[1]Benzoni, Collin-Duffresne, and Goldstein (2007), Lynch and Tan (2011) and Pastor and Stambaugh (2012) show that this conclusion can be reversed under certain conditions.

[2]In models without labor income Kim and Omberg (1996), Brennan, Schwartz and Lagnado (1997), Brandt (1999), Campbell and Viceira (1999), Balduzzi and Lynch (1999), Barberis (2000), Campbell et. al. (2001 and 2003), Wachter (2002), Liu (2007), Lettau, and Van Nieuwerburgh (2008), and Johannes, Korteweg and Polson (2014) among others, show that optimal stock market exposure varies substantially as a response to time variation in the equity risk premium.

[3]The portfolio choice literature is not limited to the papers studying time variation in the equity risk premium. For example, Munk and Sorensen (2010) and Koijen, Nijman, and Werker (2010) focus on time variation in interest rates and bond risk premia, while Brennan and Xia (2002) study the role of inflation. Chacko and Viceira (2005), Fleming, Kerby and Ostdiek (2001 and 2003) and Muir and Moreira (2017a and 2017b) consider time variation in volatility, while Buraschi, Porchia and Trojani (2010) incorporate time-varying correlations.

replicated through simple strategies that can be easily implemented by improved target date funds, in the same spirit as the optimal life-cycle strategies are replicated by currently available TDFs. Building on our initial discussion, we refer to those modified funds as Tactical Target Date Funds (hereafter TTDFs).

Our focus on the predictability driven by the VRP is motivated not only by its empirical success as a predictive factor but also by the high-frequency nature of this time variation in expected returns. More traditional predictive variables, such as CAY or the dividend-yield, capture more low frequency movements and tend to associated with bad economic conditions and/or discount rate shocks, both of which might affect households directly.[4] On the other hand, the VRP predicability is more likely driven by constraints on banks, pension funds and mutual funds (e.g. capital constraints or tracking error constraints). Such high frequency predictability is highly unlikely to be correlated with individual labor income, and individual households do not face such institutional or shareholder constraints.[5] As a result households are in a prime position to "take the other side" and exploit this premium. Furthermore, in general equilibrium households naturally own the banks and the wealth invested in the pension/mutual funds and this actually adds a further motivation for taking the other side of the VRP. If those institutional investors are forced to scale down their risky positions because of exogenous constraints, then households should be keen to offset this by increasing the risk exposure in their individual portfolios.

In that respect our paper differs from Michaelides and Zhang (2017) who incorporate stock market predictability through the dividend-yield in a life-cycle model of consumption and portfolio choice. More importantly, and differently from the previous literature on predictability, the focus of our paper is not on quantifying the welfare gains from following an optimal policy. Instead we use the output of model to design an approximate portfolio rule that can be easily implementable by an improved target date fund and thus be transparently communicated to investors. This is an important consideration since the individual investors

[4] Bad economic conditions will tend to be associated with negative labor income shocks, and discount rate shocks might reflect increased risk aversion from households.

[5] We do not have quarterly data on individual-level labor income for stockholders, but aggregate economic conditions do not fluctuate as much at a quarterly frequency, thus making it extremely unlikely that household income shocks will be correlated with the VRP factor.

are the ones who decide where to allocate their retirement savings, and several of them have limited financial literacy and might be skeptical about complex financial products.[6] Furthermore, we show that this approximate portfolio rule is able to capture a significant fraction of welfare gains implied by the optimal policy functions from the model.

Relative to an investor that assumes i.i.d. expected returns, the investor that exploits the predictability of the VRP (henceforth VRP investors) accumulates substantially more wealth by age 65, with increases in excess of 100% across a wide range of preference parameters. These implied welfare gains are also quite large, with an age-65 certainty equivalent gain of 11.5% for the baseline values of the preference parameters. In other words, the VRP investor will benefit from an 11.5% higher (risk-adjusted) consumption level at retirement than an otherwise identical investor that ignores time variation in expected returns. Furthermore these gains remain substantial when expressed in terms of life-time utility, with a life-time certainty equivalent gain of 2.0%. For comparison, also in the context of a standard model of consumption and life-cycle portfolio choice, Cocco, Gomes and Maenhout (2005) report life-time certainty equivalent losses of 2.5% from not investing in equities at all. The welfare gains that we document are even higher for investors with low risk aversion. Although they accumulate less financial wealth these investors are more willing to take advantage of return predictability and they care less about the additional consumption volatility induced by the VRP strategy.

Having documented large welfare gains from following the optimal decision rules derived in the model we turn to the main question that we wish to explore in our paper. Designing improved TDFs that are both transparent and easy to implement and yet can replicate, as much as possible, those welfare gains. Existing target date funds do not use the exact policy functions of individual households, they instead offer an approximation that can be implementable at low cost. For example, the exact policy functions imply different portfolio allocations for investors with different levels of wealth (relative to future labor income).[7]

[6] There is a growing literature documenting the low levels of financial literacy in the population at large. Lusardi and Mitchell (2014) provide an excellent survey. Guiso, Sapienza and Zingales (2008) show that trust is an important determinant of stock market participation decisions.

[7] In a similar spirit to ours, Dahlquist, Setty and Vestman (forthcoming) study simple adjustments to the portfolio rules of TDFs to take this into account.

Also, the optimal life-cycle asset allocation is actually a convex function of age as the investor approaches retirement, not a linear one. However, the approximate rule is easier to understand for investors that might have limited financial literacy, and they are the ones who decide where to allocate their retirement savings. Therefore, in the same spirit as current TDFs, we approximate the optimal asset allocations with simple linear rules that can be followed by a Tactical Target Date Fund. We estimate the best linear rule from regessions on our simulated data, where we include as explanatory factors not only age, but also the predictive factor (i.e. the volatility risk premium).[8] We further truncate the fitted linear rule by imposing fully binding short-sale constraints. It might be hard for funds taking short positions to be allowed in some pension plans, and even if that is not a concern they might be a tough sell among investors saving for retirement that have (on average) limited financial education.

We find that this simple rule generates substantial increases in age-65 wealth accumulation and in certainty equivalent welfare gains (both life-time and at retirement). In our analysis we take into account for a potential increase in transaction costs implied by the VRP strategy. Even as we consider a 0.25% decrease in expected returns due to those additional trading costs the life-time certainty equivalent gain from the TTDF versus the standard TDF is still 1%, for our baseline calibration. The expected age-65 wealth accumulation is 78% higher and the age-65 certainty equivalent gain is 3.6%. Consistent with the previous results, we find that the gains are particular higher for investors with moderate or low risk aversion. From this we can caonclude that, if the TTDFs are introduced, then those investors are the ones that would benefit the most from switching from standard TDFs into these new products.

Given that one drawback of the TTDF is that it implies significant turnover, we next consider versions of the fund were we explicitly restrict quarterly turnover to a maximum threshold. It is particularly interesting to discuss the case we set this threshold such that the average turnover of the constrained TTDF is comparable (even slightly lower) than the average turnover of the typical mutual fund (78% from Sialms, Starks and Zhang (2013)).

[8] We also explore more sophisticated rules which naturally deliver higher wealth accumulation and utility gains but, for reeasons just discussed, this one will be our baseline case.

Although the increases in expected wealth accumulation are now smaller, the turnover constraint also decreases the volatility of wealth/consumption. Therefore, even when we impose this constraint the certainty equivalent gains are almost unchanged and therefore remain very large. For the baseline parameter values the certainty equivalent gains from the TTDF are between 1% and 1.4%, depending on the assumption about the impact of the (additional) transaction costs.

We further show that different natural extensions to the proposed TTDF can lead to even larger welfare gains. Those extensions include relaxing the short-sale constraints, considering a portfolio rule where we allow the age effects to interact with the predictive factor, and extending the TTDF beyond age 65 by adding a linear portfolio rule for the retirement period also. Despite the improved results we believe that all of the above face non-trivial implementation problems relative to the simpler TTDF, which is why we only present them as extensions to our baseline case.

The paper is organized as follows. Section II outlines the theoretical life-cycle model, outlines the numerical solution algorithm and discusses the parameter choices for the calibration at a quarterly frequency level. Section III describes the data and the estimations used to calibrate the model. In Section IV we discuss the optimal portoflio strategy of the VRP investor and compare with that of i.i.d. investor. Section V discusses the design of the proposed TTDFs and in Section VI we explore different extensions. Section VII concludes.

2 The Model

Time is discrete, but contrary to previous literature we solve the model at a quarterly rather than an annual frequency. This is crucial to capture the higher-frequency predictability in expected returns documented by Bollerslev et al. (2009). Households start working life at age 20, retire at age 65, and live (potentially) up to age 100.[9] In the notation below we will use t to denote calendar time and a to denote age.

[9] Since we have a quarterly model this implies solving 324 time periods. Life-cycle models are typically solved at an annual frequency, and their inputs taken accordingly, so we discuss below how we calibrate the model to be consistent with decision making at the quarterly frequency.

2.1 Assets and Returns

In the model there are two financial assets available to the investor. The first one is a riskless asset repesenting a savings account or a short-maturity T-bill. The second is a risky asset which corresponds to a diversified stock market index. The riskless asset yields a constant gross after tax real return, R_f, while the gross real return on the risky asset is denoted by R, and its expectation is potentially time varying. The time variation in expected returns is captured by a predictive factor (f_t) and, similar to Campbell and Viceira (1999) and Pastor and Stambaugh (2012), we construct the following VAR,

$$r_{t+1} - r_f = \alpha + \beta f_t + z_{t+1}, \tag{1}$$

$$f_{t+1} = \mu + \phi(f_t - \mu) + \varepsilon_{t+1}, \tag{2}$$

where r_f and r_t denote the net risk free rate and the net stock market return, respectively. The two innovations $\{z_{t+1}, \varepsilon_{t+1}\}$ are i.i.d. Normal variables with mean equal to zero and variances σ_z^2 and σ_ε^2, respectively. The formulation allows for contemporaneous correlations between z_{t+1} and ε_{t+1}.[10]

Building on Bollerslev, Tauchen and Zhou (2009) and Bollerslev, Marrone, Xu, and Zhou (2014) we assume that the predicitive factor (f_t) is the variance risk premium (VRP_t), defined as the difference between the option-implied volatility of the stock marker (IV_t) and its realized volatility (RV_t),

$$f_t \equiv VRP_t \equiv IV_t - RV_t \tag{3}$$

We follow Bollerslev et al. (2009) in computing the two variables on the right hand side of equation (3).[11]

For comparison we will also be reporting results from a model with i.i.d. excess returns,

[10] Unlike most commonly used predictors of expected returns, the factor that we consider in this paper (the variance risk premium) is not very persistent. Nonetheless, for generality sake, in the numerical solution of the model we approximate this VAR using Floden (2008)'s variation of the Tauchen and Hussey (1991) procedure, designed to better handle the case of a very persistent AR(1) process.

[11] The details are provided in the Estimation and Calibration section.

in which case

$$r_{t+1} - r_f = \mu + z_{t+1} \tag{4}$$

. In order for the i.i.d. model to be comparable to the factor model, the first two unconditional moments of returns are set to be equal in both cases. We will also consider cases where additional transaction costs from more active trading negatively impact the expected return earned by the fund that exploits the VRP predictability. This will be implemented by adjusting appropriately the value of α in equation (1).

2.2 Preferences and Budget Constraint

The household has recursive preferences defined over consumption of a single non-durable good (C_a), as in Epstein and Zin (1989) and Weil (1990),

$$V_a = \max \left\{ (1-\beta) C_a^{1-1/\psi} + \beta \left(p_{a+1} E_a (V_{a+1}^{1-\gamma}) \right)^{\frac{1-1/\psi}{1-\gamma}} \right\}^{\frac{1}{1-1/\psi}}, \tag{5}$$

where β is the time discount factor, ψ is the elasticity of intertemporal substitution (EIS) and γ is the coefficient of relative risk aversion. The probability of surviving from age a to age $a+1$, conditional on having survived until age a is given by p_{a+1}.

At age a, the agent enters the period with invested wealth W_a and receives labor income, Y_a. Following Gomes and Michaelides (2005) we assume that an exogenous (age-dependent) fraction h_a of labor income is spent on (un-modelled) housing expenditures. Letting α_a denote the fraction of wealth invested in stock at age a, the dynamic budget constraint is

$$W_{a+1} = [\alpha_{a-1} R_t + (1-\alpha_a) R_f](W_a - C_a) + (1 - h_{a+1}) Y_{a+1} \tag{6}$$

where R_t is the return realized that period (so when $t = a$).In the baseline specification we assume binding short sales constraints on both assets, more precisely

$$\alpha_a \in [0, 1] \tag{7}$$

In practice it is expensive for households to short financial assets and relaxing these assumptions would require introducing a bankruptcy procedure in the model. In the context of the life cycle fund shorting will be cheaper, but still not costless, and this will still require making assumptions about the liquidation process in case of default. For these reasons the baseline model assumes fully binding short-selling constraints but we will also discuss results where we relax these.

2.3 Labor Income Process and Normalization

The labor income follows the standard specification in the literature (e.g. Cocco et al. (2005)), such that the labor income process before retirement is given by[12]

$$Y_a = \exp(g(a))Y_a^p U_a, \quad (8)$$

$$Y_a^p = Y_{a-1}^p N_a \quad (9)$$

where $g(a)$ is a deterministic function of age and exogenous household characteristics (education and family size), Y_a^p is a permanent component with innovation N_a, and U_a a transitory component of labor income. The two shocks, $\ln U_a$ and $\ln N_a$, are independent and identically distributed with mean $\{-0.5 \times \sigma_u^2, -0.5 \times \sigma_n^2\}$, and variances σ_u^2 and σ_n^2, respectively. We allow for correlation between the permanent earnings innovation ($\ln N_a$) and the shocks to the expected and unexpected returns (ε_{a+1} and z_{a+1}, respectively).

As also common in the literature the retirement date is exogenous ($a = K$, corresponding to age 65) and income is modelled as a deterministic function of working-time permanent income

$$Y_a = \lambda Y_K^p \text{ for } a > K \quad (10)$$

where λ is the replacement ratio of the last working period permanent component of labor income.

The unit root process for labor income is convenient because it allows the normalization of the problem by the permanent component of labor income (Y_a^p). Letting lower case letters

[12] We are assuming that the quarterly data generating process for labor income is the same as the one at the annual frequency. The calibration section discusses this in more detail.

denote the normalized variables the dynamic budget constraint becomes

$$w_{a+1} = \frac{1}{N_{a+1}}[r_{t+1}\alpha_{ia} + r_f(1-\alpha_{ia})](w_a - c_a) + (1-h_{a+1})\exp(g(a+1))U_{ia+1}. \tag{11}$$

3 Estimation and Calibration

3.1 VAR model for stock returns

The stock market data come from the Center for Research in Securities Prices (CRSP). We use the quarterly bond returns, the CPI growth rate to compute inflation, daily value-weighted cumulative returns and daily value-weighted returns of the CRSP US porftfolio index from Jan. 1st, 1990 to Dec. 31st, 2015 to construct the relevant series. The quarterly cumulative and ex-dividend return are constructed from the monthly return of the value-weighted CRSP portfolio index.

From equation (3), to construct the variance risk premium we need both the implied volatility from index options and the stock market realized volatility. The data for the quarterly implied volatility index (IV_t) is taken from the Federal Reserve Bank of St. Louis. We construct the quarterly realized volatility as in Bollerslev et al. (2009),

$$RV_t \equiv \sum_{j=1}^{n} \left[p_{t-1+\frac{j}{n}} - p_{t-1+\frac{j-1}{n}(\Delta)}\right]^2, \tag{12}$$

where RV_t is the return variation between $t-1$ and t and p_t is the natural log of the daily stock price.

Table 1 contains the descriptive statistics from the data set. The quarterly mean real free rate is 0.18% and its standard deviation is very low, and we will therefore assume it to be constant. The stock market return has a quarterly mean of 1.98% with a standard deviation equal to 7.9%. Following the life-cycle portfolio choice literature we assume an unconditional equity premium below the historical average, namely 4% at an annual frequency.

Figure 1 shows the time series variation in implied volatility (IV_t), realized volatility (RV_t) and the variance risk premium (VRP_t). Figure 1 replicates and extends essentially the original Bollerslev, Tauchen and Zhou (2009) measure. Table 2 reports the estimation results for the VAR model (1 and 2). Our quantitative estimates are largely consistent with the ones in Bollerslev et al. (2009). The factor innovation is very smooth with a standard deviation (σ_ε) of 0.007. Given these estimates, we can infer that the unconditional variance of unexpected stock market returns from

$$\sigma_z^2 = Var(r_t) - \beta^2 \sigma_f^2 \tag{13}$$

The correlation between the factor and the return innovation ($\rho_{z,\varepsilon}$) is an important parameter in determining the hedging demands. For most common predictors in the literature (e.g. dividend yield and CAY) this is a large negative number (see, for example, Campbell and Viceira (1999) and Pastor and Stambaugh (2012)). By contrast, when the predictive factor is the VRP, this correlation is estimated as slightly positive, suggesting that hedging demands are not particularly important in this context.[13]

3.2 Income process and housing expenditures

As previously discussed we consider the typical income process in the household finance literature and therefore for the most part we use the estimates in Cocco et al. (2005), which are based on the PSID. We take their estimated deterministic component of labor income ($g(a)$) and linearly interpolate in between years to derive the quarterly counterpart. Likewise we use their replacement ratio for retirement income ($\lambda = 0.68$). Cocco et al. (2005) estimate the variances of the idiosyncratic shocks around 0.1 for both σ_u and σ_n, at an annual frequency, Since we assume that the quarterly frequency model is identical to the annual frequency model it can then be shown that the transitory variance (σ_u^2) remains the same as in the annual model while the permanent variance (σ_n^2) should be divided by four.

Angerer and Lam (2009) note that the transitory correlation between stock returns and labor income shocks does not empirically affect portfolios and this is consistent with simulation results in life cycle models (Cocco, Gomes, and Maenhout (2005)). We therefore set the correlation between transitory labor income shocks and stock returns equal to zero.

[13] Indeed, if we set $\rho_{z,\varepsilon}$ equal to zero in our model the results are not significantly different. For that reason we do not explore the role of hedging demands in the paper, but those results are available upon request.

The baseline correlation between permanent labor income shocks and unexpected stock returns ($\rho_{n,z}$) is set equal to 0.15, consistent with the mean estimates in most empirical work (Campbell et. al. (2001), Davis, Kubler, and Willen (2006), Angerer and Lam (2009) and Bonaparte, Korniotis, and Kumar (2014)). We set the correlation between the innovation in the factor predicting stock returns and the permanent idiosyncratic earnings shocks ($\rho_{n,\varepsilon}$) to zero.

Finally we take the fraction of yearly labor income allocated to housing from Gomes and Michaelides (2005). This process is estimated from Panel Study Income Dynamics (PSID) and includes both rental and mortgage expenditures. As before, to obtain an equivalent quarterly process we linearly interpolate across years.

4 Optimal strategies within the model

We first document the optimal life-cycle portfolio allocations in the model with time-varying expected returns (henceforth VRP model) for a baseline value of preference parameters for the investor (henceforth VRP investor). We then document the utility gains and differences in the implied distribution of wealth accumulation at retirement relative to the case where the household follows the decision rules from the i.i.d. model. We then report some comparative statics showing how the results vary as we change the preference parameters, which indicates how the conclusions might differ across households. These results will form the basis for the next section, were we propose the tactical target date funds (TTDFs).

4.1 Optimal portfolio allocation

In the VRP model the optimal asset allocation is determined by age, wealth and the realization of the predictive factor (the variance risk premium). In Figure 2 we plot the average share invested in stocks for the VRP investor when the factor is at its unconditional mean ($\alpha_a[E(f)]$), the mean share across all realizations of the factor ($E[\alpha_a(f)]$), and the one obtained under the i.i.d. model ($E[\alpha_a^{iid}]$). In all cases wealth accumulation is being computed optimally using the appropriate policy functions. The portfolio share from the i.i.d. model follows the classical hump-shape pattern (e.g. Cocco, Gomes and Maenhout (2005)).[14] The optimal allocation of the VRP investor, for the average realization of the predictive factor ($\alpha_a[E(f)]$), shares a very similar pattern and, except for the period in which both are constrained at one, we have

$$\alpha_a[E(f)] < E[\alpha_a^{iid}] \tag{14}$$

Even though under the two scenarios the expected return on stocks is the same, Figure 2 shows that $\alpha_a[E(f)]$ is below one already before age 35 and from then onwards it is always below $E[\alpha_a^{iid}]$. The main driving force behind this result is the difference in wealth accumulation of the two investors. As we show below, the VRP investor is richer and therefore allocates a smaller fraction of her portfolio to risky assets.[15]

We next compare the optimal risky share for the average realization of the factor ($\alpha_a[E(f)]$), with the optimal average risky share across all factor realizations ($E[\alpha_a(f)]$). If the portfolio rule were a linear function of the factor the two curves should overlap exactly. However, Figure 2 shows that there is a substantial difference between the two, particularly early in life. At this early stage of the life-cycle, age below 45, we have

$$E[\alpha_a(f)] < \alpha_a[E(f)] \text{ for } a < 45 \tag{15}$$

This result arises from a combination of the short-selling constraints and the fact that $\alpha_a[E(f)]$ is (much) closer to one than to zero. Given the high average allocation to stocks early in life, for realizations of the factor above its unconditional mean the portfolio rules are almost always constrained at one. On the other hand, for lower realizations of the predictive factor the optimal allocation is "free" to decrease, hence it is lower than $\alpha_a[E(f)]$. As a result, optimal allocation of the VRP investor is sometimes far below $\alpha_a[E(f)]$ and never exceeds it by much.[16] Building on this intuition, it is not surprising to find that the sign of inequality flips once the portfolio allocation at the mean factor realization ($\alpha_a[E(f)]$)

[14] The increasing pattern early in life is barely noticeable because under our calibration the average optimal share at young ages is (already) close to one.

[15] The two policy allocations also differ because the policy rules from the VRP model take into account the hedging demands, but that effect is quantitatively much less important.

[16] It is similar to averaging a truncated distribution where the tructation is mostly binding at the upper limit.

falls below 50%, which takes place around age 45. Now the more binding constraint is the short-selling constraint on stocks so we have:

$$E[\alpha_a(f)] > \alpha_a[E(f)] \text{ for } a > 45 \tag{16}$$

This comparison suggests that the welfare gains from the VRP model are likely to be much higher if we relax the short-selling constrains, which motivates our discussion of this particular extension in Section 6.

Combining the previous two inequalities (14) and (15) it is easy to see that, until age 45, we have:

$$E[\alpha_a(f)] < E[\alpha_a^{iid}] \tag{17}$$

namely that the average portfolio allocation in the VRP model ($E[\alpha_a(f)]$) will be much lower than the one in the i.i.d. model ($E[\alpha_a^{iid}]$), and the intuition follows from the previous discussions. In fact, even after age 45, when (15) is replaced by (16), we see that, although the difference between the optimal allocation of the VRP and i.i.d. investors decreases, equation (17) still holds: inequality (14) dominates inequality (15).

4.2 Portfolio returns, wealth accumulation and utility gains

In this section we first discuss the differences in expected returns between the VRP and i.i.d. investors and then the implied differences in wealth accumulation at retirement and certainty equivalent consumption levels. In these calculations we ignore transaction costs merely to avoid repetition, as we will naturally consider them in the next section when we discuss the implementation of these portfolio rules in the context of the improved target-date funds.

In Figure 3 we plot the (annualized) average expected portfolio returns over the life-cycle

$$E(R_{t+1}^P) = \alpha_a E_t[R_{t+1}] + (1 - \alpha_a)R_f,\ a = 1, ..., T \tag{18}$$

computed by averaging at each age across all simulations.

Since we are averaging across all possible realizations of the factor, for a constant portfolio allocation ($\overline{\alpha}$) this would be a flat line. For example, if $\overline{\alpha} = 1$, this would be equal to the average equity portfolio return, regardless of age. In the i.i.d. model this line essentially inherits the properties of the optimal $\{\alpha_a\}_{a=1}^{T}$. The (annualized) expected portfolio return is around 5% early in life, increases slightly in the first years and then decays gradually as the investor approaches retirement and thus shifts towards a more conservative portfolio. In the VRP model the same average life-cycle pattern is present but now, since the household increases (decreases) α_a when the expected risk premium is high (low), the line is shifted upwards. As a result, even though as shown in Figure 2 the VRP investor has on average a lower exposure to stocks than the i.i.d. investor, her expected return is actually higher.

The vertical difference between the two lines gives us a graphical representation of the additional expected excess return that is actually earned by the VRP investor, and to facilitate the exposition we also plot it as a separate line in the figure. From age 37 onwards this difference increases monotonically, as the lower average equity share makes the short-selling constraint less binding and thus the VRP investor is more able to exploit time-variation in the risk premium. As the two agents reach retirement, the difference in expected returns is almost 4%. This difference is therefore at its maximum exactly when these investors have the highest wealth accumulation.

In Figure 4 we compare the average wealth accumulation of the i.i.d. and the VRP investors over time. In order to facilitate the interpretation wealth is scaled by (annualized) age-20 income, as common in the life-cycle literature. Consistent with the differences in expected returns documented in figure 3, the wealth of the VRP investor grows at a much faster rate than that of the i.i.d. investor, and as a result she accumulates 140% more wealth by retirement age: 17.0 versus 7.1. The implied welfare implications for retirees are significant, with an age-65 certainty equivalent gain of 11.5%. This is computed as the difference in the certainty equivalent consumption levels at retirement for the VRP investor and for an investor that ignores predictability. In other words, the VRP investor will expect a 11.5% higher risk-adjusted consumption level per year, from age 65 onwards.

It is important to point out that the VRP investor also has a 13% higher average yearly consumption *before* retirement. So the higher retirement consumption is not being financed by lower savings, quite the contrary. The strategy allows the investor to obtain a higher level of expenditure throughout her life. Reflecting this, the life-time equivalent gain (i.e. the certainty equivalent gain at age 20) is 2%. This number is very large when compared with typical welfare cost calculations. For example, Cocco, Gomes and Maenhout (2005) report life-time certainty equivalent losses of 1.5% for investors that ignore labor income altogether when chosing their optimal portfolios, and life-time certainty equivalent losses of 2.5% from not investing in equities at all.[17] The fact that the welfare gains from exploiting the time-variation in expected returns based on the VRP are quite similar in magnitude to these confirms their economic importance.

4.3 Comparative statics

The results presented so far were obtained under our baseline calibration of the preference parameters. In order to have a more complete understanding of how the market timing strategy might impact different households we now consider alternative calibrations. In Table 3 we report the average risky share at different ages, age-65 wealth accumulation and corresponding certainty equivalent gain, for different values of risk aversion (2, 5 and 10).[18]

In Table 3, Panel A we report the average allocation to stocks at different ages over the working part of the life cycle and the standard deviation of the share of wealth in stocks. As we increase risk aversion the average allocation to stocks naturally falls. This result is less pronounced early in life, when the allocation for a large range of realizations of the predictive factor is constrained at 100%, as is the case for $\gamma = 2$ and $\gamma = 5$, hence giving an unconditional average asset allocation at around 70%. But the pattern becomes quite clear as the investor ages. The cross sectional standard deviation of the share of wealth in stocks is higher for lower risk aversion coefficients: 44% for the investor with risk aversion of 2, versus 40% and 37% respectively for risk aversion of 5 and 10. This reveals that the less risk-averse investors are more willing to explore time-variation in the risk premium.[19] Intuitively they care less about the additional portfolio volatility (and hence consumption

[17] These are computed in the context of a standard life-cycle model of consumption and portfolio decisions, very similar to ours, with the same calibration for returns and the same coefficient of relative risk aversion.

[18] Comparative statics for the other preference parameters are available upon request.

[19] The standard deviation naturally also reflects fluctuations in the portfolio share due to changes in wealth accumulation and age effects, just as in the i.i.d. model.

volatility) that this activity generates.

In Table 3, Panel B we compare the VRP investor with an otherwise identical i.i.d. investor in terms of age-65 wealth accumulation, pre-retirement consumption and certainty equivalent gains. The first row documents that the increase in age-65 wealth is higher for the less risk-averse investors: 167%, 140% and 113% for γ of 2, 5, and 10, respectively. This is driven by two effects. First, as documented in Panel A, the less risk averse investor is more eager to exploit time variation in expected returns, reflected in the higher standard deviation of the portfolio share during working life. Second, as shown in row 3 of Panel B, there is a differential increase in consumption before age 65, with the more risk-averse investors increasing their expenditures more during this period, and thus saving less for retirement.

The wealth accumulation results are also affected by the presence of the short-selling constraints. These constraints limit the ability to exploit time variation in expected returns but their impact is more complex since they affect different investors differently, depending on their average portfolio allocation. Those with an average allocation of 50% are less affected than those with an average allocation of 75% (25%), for example. The second investor is less able to exploit states with high (low) expected returns.[20] From this intuition we see that this effect works particularly against the investors with both low and high risk aversion. Therefore, it is not clear ex-ante that less risk-averse investor will experience a more substantial increase in wealth accumulation. This result can only be obtained from solving the calibrated model as we have done.

In Panel B we also report the welfare gains at age 20 and at age 65. When we measure the age-65 gains in utility terms, in addition to the average wealth accumulation effect, we are also taking into account that the market timing strategy also implies an increase in the standard deviation of age-65 wealth. As shown in the second row of Panel B, these increases are quite large.[21] This highlights why it is important to measure the benefits in terms of

[20] The investors with the 75% (25%) average allocation can partially compensate for this by being able to fully exploit states with even lower (higher) expected returns but, by definition, those states have low probability.

[21] We are reporting the increases as a percentage, since we believe it makes it easier to compare numbers. If we reported the change in levels they would be monotonically increasing in risk aversion. Also, as discussed when presenting the wealth accumulation results, the presence of the short-selling constraints induces non-

certainty equivalent gains, otherwise we would be over-estimating the benefit of the market timing rules.[22]

The previous result also explains why the differences in age-65 utility gains as we change risk aversion are even more significant economically than the ones obtained for wealth accumulation: the certainty equivalent gain of the investor with a risk aversion of 2 is 46.0%, versus 11.5% for a risk aversion of 5, and 4.6% for a risk aversion of 10. The VRP strategy implies a large increase in the standard deviation of (age-65) wealth, and by definition the less risk-averse investor is the one caring less about this additional risk. As a result, the utility of the investor of risk aversion of 2 is less affected by the added risk. Combining this with the fact that she also accumulates more wealth it is not surprising to find that her gain, when measured in certainty equivalent consumption units, is much larger than the one of the other investors.

In the last columns of Panel B we report the life-time certainty equivalent gains, and these share the same patterns as those at age-65. The welfare gain for the investor with risk aversion of 2 is 2.6% confirming that the less risk-averse investors stand to benefit significantly from the market timing rule.

5 Tactical Target-date Funds

In the previous section we have shown that exploiting the equity premium predictability from the volatility risk premium generates significantly higher expected wealth accumulation at retirement and leads to large utility gains. However, those gains were computed for an investor using the optimal policy functions from the model, which is not a feasible solution for a mutual fund. Target date funds do not use the exact policy functions of individual households, they instead offer an approximation that can be implementable at low cost. This approach benefits from the further advantage that such a simpler strategy can be more easily communicated to investors that might have limited financial literacy, and are the ones who

monotonic patterns, as both the very high and the very low risk averse investors are the ones that become more constrained.

[22] Later on we will present results for constrained versions of the market timing rule, for which the standard deviations of wealth are much lower.

decide where to allocate their retirement savings.

The current practice therefore is for the vast majority of target-date funds (TDFs) to approximate the optimal life-cycle risky share using a linear function of age. This is an approximation to the typical optimal solution for the i.i.d. model which follows a hump shape pattern early in life, although not very pronounced for low levels of risk aversion, and has a convex shape later on as the investor approaches retirement. However, as the exact patterns of optimal policy will vary across individuals based on their preferences and other important factors (e.g. labor income profile and wealth accumulation), the linear function is thus viewed as simple to explain and a reasonable approximation to an heterogeneous set of optimal life-cycle profiles.[23]

In the same spirit, in the baseline specification we derive a relatively straightforward portfolio rule that can be implemented by an improved target date fund (the TTDF) and which will aim to capture a large fraction of the welfare gains previously described. More precisely we now derive optimal policy rules that consist of linear functions of age and of the predictive factor. If we design more complicated rules we could potentially increase the certainty equivalent gains, and in fact we also explore some alternative portfolio rules along these lines. Finally, in this section, both for the i.i.d. and for the VRP cases, we further constrain the estimated portfolio rules by forcing them to satisfy the short-selling constraints. Later on we discuss the results obtained when we relax this constraint.

5.1 Optimal target-date funds

5.1.1 Tactical TDF with the VRP as a regression covariate (TTDF)

The simplest extension of the traditional TDF portfolio that incorporates the predictability channel is obtained by adding the predictive factor as an additional explanatory variable in a linear regression. More precisely, we use the simulated output from the model to estimate

$$\alpha_{iat} = \theta_0 + \theta_1 * a + \theta_2 * f_t + \varepsilon_{iat}. \quad (19)$$

[23] Dahlquist, Setty and Vestman (forthcoming) explore improved portfolio rules that take into account household-specific characteristics.

Relative to the optimal simulated profiles this regression is quite restrictive as, in addition to linearity, it implies that both the regression coefficient on age (θ_1) and the intercept (θ_0) are the same regardless of the realization of the factor state. However, as previously argued, this is simple to implement and easier to explain to investors..

Table 4, Panel A and Figure 5 report the regression results from these rules for the baseline case of relative risk aversion equal to 5 and, for comparison, the results for the i.i.d. model.[24] Panel B reports the fitted linear rules for other values of risk aversion (2 and 10). These would correspond to three different TTDFs, each targeted to investors with different levels of risk aversion.

The life-cycle asset allocations for both the i.i.d. and the VRP baseline model are reasonably well captured by a linear regression rule. Despite the higher complexity of the optimal portfolio rules in the VRP case, the R-squared of the fitted linear regression is actually higher: 74% versus 45%. This is due to the lower implied average allocation to stocks, as already documented in Figure 2, which makes the short-selling constraints less binding. In the regression specification age is expressed in quarters starting for quarter 1, as in the model. Therefore, the rule age pattern for the i.i.d. case is slightly steeper than the popular "100-age" rule followed by several existing target-date funds, but not far away from it. Similarly, the average age pattern of the VRP rule is slightly flatter than the 100-age rule but, likewise, not very different from it. Of course under the VRP rule (equation (19)) the allocation also changes with the predictive factor. For example, for sufficiently high (or sufficiently low) values of this factor, the short-selling constraints can become binding. Later on, when evaluating these strategies, we discuss their implied turnover.

In the last two columns of Table 4 we report the regression results for different values of relative risk aversion. As risk aversion decreases the coefficient on the predictive factor increases (in absolute value), consistent with the discussion in the previous section. The less risk averse investor is more willing to take advantage of time variation in expected returns. However, as also previously discussed, given that the less risk averse investor has an average portfolio allocation that is much closer to 1, her ability to actually follow the optimal

[24] These are regressions on data simulated from the model so the t-statistics are all extremely high almost by definition, and therefore are omitted from the table.

market timing strategy is more limited by the presence of the short-selling constraints. This is reflected in the significantly lower regression R^2: 58 percent versus 74 (73) percent for relative risk aversion equal to 5 (10).

5.1.2 Tactical TDF conditioning on the VRP (TTDF2)

As previously discussed, the portfolio rule based on equation (19) is very straightforward but quite restrictive. Therefore, we also consider an alternative formulation where we fit the simulated shares of wealth in stocks on age using separate regresssions conditional on the different realizations of the predictive factor. So, we run the following series of regressions for each f_j in our discretization grid

$$\alpha_{iat} = I_{f_t=f_j}\theta_0^j + I_{f_t=f_j} * \theta_1^j * a + \varepsilon_{iat}^j, \text{ for each } f_j \tag{20}$$

where $I_{f_t=f_j}$ equals to 1 if $f_t = f_j$ and equals to 0 otherwise.

The results are shown in Table 5 and Figure 6. Panel A of Table 5 reports, for the baseline case of risk aversion equal to 5, the regression results for three different values of f_j : mean and plus and minus two standard deviations.[25] Panels B and C report the same results for risk aversions of 2 and 10, respectively.

As we can see, realization of the factor at plus (minus) two standard deviations away from the mean already imply a 100% (0%) allocation to stocks regardless of age. This pattern is not captured by the more restrictive TTDF rule (equation (19)) and is reminiscent of the Brennan, Schwartz and Lagnado (1997) results of a bang-bang solution with the intermediate cases closer to the mean having a pronounced age effect due to the presence of undiversifiable labor income.

5.2 Utility gains

Having identified a feasible portfolio rule for the TTDF we now proceed to compute the corresponding certainty-equivalent utility gains. In these calculations, as previously mentioned,

[25] As before, we again include the results for the i.i.d. investor for comparison.

we also take into account a potential increase in transaction costs implied by the market timing strategy. More specifically, we consider that the TTDF might face an effectively lower expected equity return as a result of these costs. We then report the wealth accumulation at age 65 and certainty equivalent gains from investing in the TTDFs relative to the standard TDF that ignores the market timing information provided by the realization of the factor. Results are shown for different values of risk aversion and for different assumptions about the *additional* transaction costs (tc) faced by the former.[26]

In both cases, TTDF and standard TDF we assume the same asset allocation rules at retirement, more precisely we assume that the investor ignores predictability from age 65 onwards. In other words we are measuring the gains from changing the portfolio rule in TDF only. The gains would naturally be larger if we also allowed the investor to exploit time-variation in the risk premium during retirement as well, and we present results for this case in one of our extensions below. Finally we assume that each investor can identify the fund that matches her level of risk aversion, both for the TTDFs and the standard TDF.

5.2.1 Tactical Target Date Fund 2 (TTDF2)

It is useful to start the discussion by computing the wealth and welfare changes when the more sophisticated TTDF2 rule (equations (20)) is used. This is the rule where the regressions are performed conditional on the factor realization, implying that the age effects are different across factor realizations. The results are reported in Table 6.

Comparing the results in Table 6 with those in Table 3 we see that, with the TTDF2 rule and $tc = 0$, we can replicate a very large fraction of the gains from the VRP model. The increases in wealth accumulation at age 65 are very similar. For the three different values of risk aversion these are now 167%, 140% and 113%. versus 164%, 128% and 123% in Table 3. There is a noticeable increase in the standard deviations of age-65 wealth, but the age-65 utility gains are still 6.9% for the baseline risk aversion of 5, and increase (decrease) to 31.1% (3.2%) for risk aversion of 2 (10). These numbers are about 2/3 of those obtained under

[26] The standard TDF will also face transaction costs but in our simulations we only explcitly introduce them for the enchanced fund, which is why we view them as additional costs, over and above those already faced by the standard TDF.

the optimal strategies. When looking at life-time utility the welfare gains are even closer to those in table 3, ranging from 1.2% to 2.2% for the different investors. For the reasons that we previously discussed we do not view this rule as a very practical proposition for a TDF. However, these results suggest that individuals that would be willing to invest in such funds, for example those with high financial literacy, would potentially obtain very large certainty equivalent gains.

5.2.2 Tactical Target Date Fund (TTDF)

We now study the results for the simpler TTDF rule (equation (19)). These are shown in Table 7, for different values of risk aversion (γ) and different values of the additional transaction costs (tc).

We first consider the case with $tc = 0.0$, which is directly comparable with the results in the previous section. For the investors with risk aversion of 5 and 10 both the increases in age-65 wealth accumulation and the certainty equivalent are not far below those in table 3. For example, for risk aversion of 5 the increase in age-65 wealth is 106% and the life-time certainty equivalent gain is 1.4%, which correspond to 75% and 70% of the corresponding values in Table 3, respectively. From this we see that, when the two funds face the same level of transaction costs, the simple rule is still able to capture a large fraction of the gains that were previously identified. This is even more remarkable if we recall that, in this analysis, we are assuming that the investor does not exploit the predictability in expected returns at retirement. More importantly, the welfare gains remain economically large even as we introduce the additional costs. Even when we consider a 25 basis points increase in costs age-65 wealth accumulation is still 78% higher under the TTDF with an associated life-time certainty equivalent consumption gain of 1.0%, and an age-65 certainty-equivalent gain of 3.6%.

For the investor with risk aversion of 2 the increase in age-65 wealth accumulation with $tc = 0$ is 95% versus the 167% obtained in Table 3. Likewise the certainty equivalent gains are also about half of the one obtained in the previous sections. The explanation for this comes from the regression results discussed in table 4, where we found that the TDF rule is less accurate for the investor with risk aversion of 2. In fact the life-time certainty equivalents are now marginally higher for the investor with risk aversion 5. The age-65 welfare gains are still much larger for the less risk-averse investor, as high as 18.8% for the $tc = 0$ case. This occurs because, as previously shown in table 3, these investors consume a smaller fraction of the additional wealth accumulation early in life, and thus save a larger fraction of it for retirement.

One implication of these results is that it would be particularly beneficial to introduce the TTDFs in pension plans with investors with low or moderate risk aversion (2 or 5). Along these lines, if such funds are offered in parallel with standard target date funds, those investors are the ones that would benefit the most from switching away from the conventional product.

5.3 Introducing turnover restrictions

One potential concern with the TTDFs is that its implementation might imply very high portfolio turnover. The average (annualized) portfolio turnover of the standard TDF (i.e. the one that replicates the optimal allocation of the i.i.d. investor) is 23%. For the TTDF average turnover rises to 213% indicating that tactical asset allocation implies a more volatile asset allocation behavior over the life cycle. By comparison, the average turnover of the typical mutual fund is 78% (see Sialms, Starks and Zhang (2013)).

In the previous section we included in our analysis additional transaction costs that this high turnover might generate. In this section we follow a more direct approach where we explicitly restrict the turnover of the fund. The restriction limits the optimal rebalancing of portfolio share to a maximum threshold (k). More precisely, the portfolio rule is subject to the additional constraint

$$\alpha_a = \begin{cases} \alpha_{a-1} + k & if\ \alpha_a^* > \alpha_{a-1} + k \\ \alpha_a^* & if\ |\alpha_a^* - \alpha_{a-1}| < k \\ \alpha_{a-1} - k & if\ \alpha_a^* < \alpha_{a-1} - k \end{cases} \tag{21}$$

where α_a^* is the optimal allocation in the absence of the constraint.

This constraint has the advantage of being very straightforward to implement and in our analysis we consider two thresholds, $k = 25\%$ and $k = 15\%$. We impose equation (21) directly in the simulations, instead solving the corresponding dynamic programming problem, for two reasons. First, solving the optimization problem with the constraint would require introducing an additional state variable (the lagged portfolio share, α_{t-1}). Second, even though the optimal policy function would by definition satisfy constraint (21), that does not guarantee that the corresponding fitted linear rule estimate from the simulated data would as well.[27] The results are shown in Table 8, for the case of an investor with risk aversion of 5.

With a maximum rebalancing limit of 25% the average turnover of the fund falls almost by half, to 107%. When the limit is even stricter, 15%, the average turnover is now only 69%, which is now even below that of the typical mutual fund (78% as mentioned above). Crucially, for the same the level of transaction costs as in the version without turnover restrictions (Table 7), the life-time certainty equivalent gains are identical in all cases, up to the first decimal. And in fact, since the turnover numbers are now comparable to those of existing funds we those transaction costs are likely to be lower. These results can be understood by looking at both the average wealth accumulation at age 65 and its standard deviation. The constraints do limit the fund's ability to exploit time-variation in the risk premium and this is reflected in lower expected wealth accumulation. For example, for $tc = 0$ the expected increase in age-65 wealth accumulation was 106% in the absence of the turnover constraints, but falls to 46% and 31% for $k = 25\%$ and $k = 15\%$, respectively. However, this is accompanied by an equally significant reduction in the standard deviation of (age-65) wealth: from 129% to 41% and 21%. This explains why, although the certainty equivalent gains are smaller when we introduce the constraints, the difference is negligible.

One case is particularly interesting, the one with $k = 15\%$ and $tc = 0.25\%$. Here the standard deviation of age-65 wealth is only 3% higher than in the standard TDF, yet there is a 15% difference in expected wealth accumulation. So for a very similar level of ex-ante risk the investor is obtaining a noticeable difference in average expected return. This is reflected

[27] Any mis-specification of the optimal policy functions will only lead us to under-estimate the utility gains since the constraint is more binding for the TTDF than the standard TDF.

in a certainty equivalent gain of 0.5% at age 65, and in fact an even larger life-time CE gain: 1.0%.[28]

The results in the previous sections revealed that the investor with risk aversion of 2 would benefit as much from the TTDF, as these funds were found to be particularly beneficial for investor with low and moderate risk aversion. In Table 9 we investigate the gains for this investor are affected by the introduction of the portfolio rebalancing constraint (equation (21)). The main conclusion from these results is that, just like for the investor with risk aversion of 5, the welfare gains for this investor are also unaffected by the turnover restrictions. Likewise, these restrictions are quite effective at limiting turnover, and when we set $k = 15\%$ the implied portfolio rebalancing is in line with that of standard mutual funds.

Overall, the results in Tables 8 and 9 confirm that it is possible to design a relative simple rule that exploits the risk premium predictability obtained from the VRP, requires standard levels of turnover, and is able to generate economically large welfare gains, for investors that are not very risk averse.

6 Extensions

6.1 Relaxing the short-selling constraints

As shown in Figure 6, the optimal portfolio allocation implied by the VRP strategy is sometimes constrained at either 100% or 0%. These results suggest that the utility gains from the VRP strategy are likely to be higher if we relax the short-selling or borrowing constraints. In the life-cycle asset allocation literature it is common to impose fully binding short-selling and borrowing constraints since it is particularly hard or expensive for retail investors to engage in unsecured borrowing or short-selling. However, in our case the strategy will be implemented by a mutual fund and hence it should be much cheaper and feasible to take both borrowing and short positions. Nevertheless we have considered fully binding short-selling constraints as our baseline case. We did so because a mutual fund that takes

[28] Again, for comparison, in Cocco, Gomes and Maenhout (2005) the life-time CE losses from not investing in equities at all are between 0.9% and 4.0%, for different parameter values.

leveraged positions might not be regarded as an acceptable choice by some pension plan providers.

In this section we investigate the case in which the TTDF can increase its allocation to stocks as far as 200% through borrowing at the same riskless rate, that is:

$$\alpha_a \in [0, 2] \tag{22}$$

For the range of parameter values that we consider the upper bound on this constraint becomes essentially non-binding since the optimal allocation rarely exceeds 200%. We could potentially also relax the short-selling constraint on the risky asset and the welfare gains would be even higher, but that particular constraint is less binding given that the average allocation to stocks is above 50%. Furthermore, short-selling the aggregate stock market is typically harder and more expensive to implement than borrowing to invest in stocks.

In the i.i.d. model the household borrows to invest in the stock market early in life and then the pronounced life cycle effect lowering the share of wealth in stocks takes over. We use this rule to construct the TDF for the i.i.d. model (the strategic asset allocation benchmark). In this model stock market turnover now rises to 113% relative to 23% in the benchmark analyzed earlier. We follow a similar strategy for the TTDF. Table 10 reports the differences in wealth accumulation and CE gains from taking advantage of the TTDF when we relax the short-selling constraint on the riskless asset for both funds.[29]

Comparing these results with those in Tables 7 and 8, where short-selling was completely ruled out, we find significant increases in certainty equivalent gains. Without any turnover restrictions (columns II to IV) the welfare gains more than double in size, increasing from 1.4% (1.0%) to 2.5% (1.9%) for $tc = 0$ (0.25%). The relaxing of the short selling constraint significantly increases the TTDF's ability to exploit the time-variation in the expected risk premium.

One potential concern here is that this strategy implies significantly higher portfolio turnover. In fact, we see that average fund turnover is now 310% as opposed to 213% for

[29] We maintain all other assumptions as in the baseline case, namely relative risk aversion of 5. Results for other values of risk aversion are available upon request.

the case with fully binding short-selling constraints. In the rest of Table 10, we therefore repeat the analysis and impose an exogenous constraint on trading (equation (21)). As we introduce the tigher version of constraint ($k = 15\%$) portfolio turnover drops significantly, to around 74%. Nevertheless, and consistent with the results in the previous sections, we find that the welfare gains are barely affected, and in fact they remain the same up to the first decimal case. These results show that, by relaxing the short-selling constraint on the riskless asset the TTDF can generate life-time certainty equivalent gains of 1.9% to 2.5%, for reasonable levels of turnover.

6.2 Adding VRP strategies during retirement

In the previous section the investor only exploited time variation in expected returns before retirement, through the TTDF. The goal was to isolate the role of the TTDF and thus show how introducing these market timing strategies in a target date fund alone could improve welfare. In this section we consider the benefits of trying to capture the VRP strategy throughout the life-cycle. For this purpose we consider a combination of the simple TTDF with an otherwise equally designed fund for the retirement period. More precisely, we also run the regression given by equation (19) for age greater than 65. From this we obtain a linear portfolio rule for the retirement period which complements the TTDF.

The results are shown in Table 11, for the baseline case of risk aversion 5 and with turnover restrictions to keep trading volume consistent with that of typical mutual funds. As we would have expected the welfare gains at retirement are now much larger. For the tighter turnover restrictions the age-65 certainty equivalent gains are between 3.4% and 4.9% which compares with 0.5% to 1.9% in table 8. The increase in life-time utility is naturally smaller but still non-trivial: from $1.0\% - 1.4\%$ in the previous section to $1.3\% - 1.6\%$ now.

7 Conclusion

We have shown how TTDFs can be designed in a parsimonious way and can deliver substantial welfare gains by combining the long term strategic asset allocation perspective of a life cycle investor and the short term market information that gives rise to tactical asset allocation. In unreported experiments we have extended the analysis to out of sample results (Lan (2015)), a wider set of preference parameter configurations and different models of investor behavior during retirement. We also think that there are potentially larger welfare gains that could be even obtained by allowing for a negative position in the stock market. Looking further into the design and commercialization of the proposed TTDFs, and the potential complications, that may arise in such implementations can be an interesting topic for future research.

Table 1: Descriptive Statistics for Predictive Variables

Table 1 presents descriptive statistics of quarterly data from 1990Q1 to 2016Q3: r denotes the real return on the S&P 500 index (deflating using the consumer price index (CPI)), IV denotes the quarterly "model free" implied variance or VIX index, and RV is the quarterly "model free" realized variance. Inflation (π) is derived from CPI and r_f is the real 90-day T-bill rate. These two series and the S&P 500 index are from the Center for Research in Security Prices (CRSP).

Panel A. Summary Statistics

1990Q1 –2016Q3	r	IV	RV	$IV - RV$	π	r_f
Mean (%)	1.98	1.11	0.62	0.49	0.06	0.16
SD (%)	7.84	0.94	0.98	0.75	0.08	0.09
Kurtosis	3.24	8.16	54.23	31.83	9.64	5
Skewness	-0.4	2.25	6.45	-3.24	-1.39	0.32
AR(1)	0.0	0.41	0.47	-0.17	0.88	0.09

Panel B. Correlation Matrix

1990Q1 –2016Q3	r	IV	RV	$IV - RV$	π	r_f
r	1.00	-0.52	-0.42	-0.096	-0.11	0.096
IV	–	1.00	0.7	0.34	-0.18	0.12
RV	–	–	1.00	-0.43	-0.46	0.3
$IV - RV$	–	–	–	1.00	0.38	-0.24
π	–	–	–	–	1.00	-0.76
r_f	–	–	–	–	–	1.00

Table 2: Predictive Regressions

Table 2 presents predictive regressions based on quarterly data from the first quarter of 1990 to the third quarter of 2016. The parameters related to the predictive regression using VRP as a predictor are estimated from the following restricted VAR:

$$\begin{bmatrix} VRP_{t+1} \\ r_{t+1} - r_f \end{bmatrix} = \begin{bmatrix} Const \\ \alpha \end{bmatrix} + \begin{bmatrix} \phi & 0 \\ \beta & 0 \end{bmatrix} \begin{bmatrix} VRP_t \\ r_t - r_f \end{bmatrix} + \begin{bmatrix} \varepsilon_{t+1} \\ z_{t+1} \end{bmatrix}$$

Newey-West t-statistics are reported in parentheses (α is set to zero).

1990Q1 –2016Q3	VRP
Constant	0.0058 (6.72)
α	0.0
β	3.6 (4.48)
ϕ	-0.18 (-1.84)
$\rho_{z,\varepsilon}$	-0.04
σ_ε	0.0074
σ_z	0.0746
σ_r	0.079
Adj. R^2 (%)	15

Table 3: VRP relative to i.i.d. model and effect of preference heterogeneity.
Table 3 presents percentage changes in the share of wealth in stocks between VRP and i.i.d. model. Percentage changes reported.

Panel A: Effect of Preference Heterogeneity			
γ	2	5	10
ψ	0.5	0.5	0.5
δ	0.9875	0.9875	0.9875
α_{20-29}	70	71	48
α_{30-39}	73	64	36
α_{40-49}	72	51	33
α_{50-65}	72	45	20
$std(\alpha_{20-65})$	44	40	37
Panel B: Relative to I.I.D. Model			
W_{65} (% inc.)	167	140	113
$Std(W_{65})$ (% inc.)	82	115	81
C_{20-65} (% inc.)	9.4	12.7	21.7
Age-21 Static CE Gain	2.6	2.0	1.3
Age-21 CE Gain	2.6	2.0	1.3
Age-50 Static CE Gain	11	4.6	1.8
Age-50 CE Gain	30.0	8.6	3.6
Age-65 Static CE Gain	8.3	3.4	1.4
Age-65 CE Gain	46.0	11.5	4.6

Table 4: TDF with constant age effects across risk aversion parameters

Table 4 presents the regression of simulated portfolios on age and factor realizations across different relative risk aversion coefficients (2, 5, 10).

	VRP	$i.i.d.$	$\gamma = 2$	$\gamma = 10$
Constant	0.51	1.06	0.46	0.26
Age	-0.00191	-0.00308	-0.000312	-0.00128
Factor	45.6		45.1	43.0
R^2	74%	45%	58%	73%

Table 5: Age regressions conditional on factor realizations

Table 5 presents the regression of simulated portfolios on age conditional on each factor realization, that is, age coefficients are different across factors. The experiments are shown for different relative risk aversion coefficients (2, 5, 10).

Panel A ($\gamma = 5$)				
	VRP	VRP	VRP	$i.i.d.$
Factor	2 s.d. above mean	Mean	2 s.d. below mean	
Constant	1.00	1.06	0.0	1.06
Age	0.00	-0.0046	0.0	-0.00308
R^2	0%	79%	0%	45%
Panel B ($\gamma = 2$)				
Factor	2 s.d. above mean	Mean	2 s.d. below mean	$i.i.d.$
Constant	0.96	0.96	0.0	0.89
Age	0.0003	0.0002	0.0	0.0008
R^2	3%	1%	0%	7%
Panel C ($\gamma = 10$)				
Factor	2 s.d. above mean	Mean	2 s.d. below mean	$i.i.d.$
Constant	1.0	0.43	0.0	0.45
Age	0.0	-0.002	0.0	-0.002
R^2	0%	41%	0%	39%

Table 6: TTDF conditioning on Factor (TTDF2)

Table 6 presents results from comparing the TTDF2 with the standard TDF for different relative risk aversion coefficients and additional transaction costs from trading the TTDF2. In the standard TDF the portfolio allocation rule is a linear function of age only. Under the TTDF2 the portfolio allocation also depends on the variance risk premium (VRP), by considering different linear functions of the age for each realization of the VRP. The results are reported in percentages.

γ	2	2	2	5	5	5	10	10	10
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	164	140	107	142	128	109	129	123	113
$Std(W_{65})$ (% inc.)	191	162	121	186	170	147	257	248	234
Age-21 Static CE	2.2	1.9	1.5	1.7	1.5	1.3	1.2	1.1	1.0
Age-21 CE Gain	2.2	1.9	1.5	1.7	1.5	1.3	1.2	1.1	1.0
Age-50 Static CE	8.3	7.0	5.1	3.1	2.8	2.3	1.1	1.0	0.8
Age-50 CE Gain	25.7	22.3	17.5	7.2	6.5	5.5	2.9	2.7	2.4
Age-65 Static CE	0.26	−0.45	−1.31	0.12	−0.15	−0.51	0.05	−0.04	−0.17
Age-65 CE Gain	31.1	25.9	18.9	6.9	6.1	4.9	3.2	3.0	2.7

Table 7: TDF with factor as regressor (TTDF)

Table 7 presents results from comparing the TTDF with the standard TDF for different relative risk aversion coefficients and additional transaction costs from trading the TTDF. In the standard TDF the portfolio allocation rule is a linear function of age only. Under the TTDF the portfolio allocation also depends on the variance risk premium (VRP), which enters as an additional variable in the linear regression. The results are reported in percentages.

γ	2	2	2	5	5	5	10	10	10
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	95	77	52	106	95	78	107	101	92
$Std(W_{65})$ (% inc.)	80	60	34	129	115	95	206	197	183
Age-21 Static CE	1.4	1.1	0.8	1.4	1.2	1.0	1.1	1.0	0.8
Age-21 CE Gain	1.4	1.1	0.8	1.4	1.2	1.0	1.1	1.0	0.8
Age-50 Static CE	5.0	3.9	2.5	2.6	2.3	1.8	1.0	0.9	0.7
Age-50 CE Gain	16.4	13.6	9.7	5.9	5.2	4.3	2.6	2.4	2.1
Age-65 Static CE	0.15	−0.43	−1.12	0.1	−0.15	−0.47	0.04	−0.04	−0.16
Age-65 CE Gain	18.8	14.8	9.4	5.5	4.7	3.6	2.8	2.6	2.2

Table 8: Results with turnover restrictions when $\gamma = 5$

Table 8 presents results from comparing the TTDF with the standard TDF for different rebalancing restrictions and transaction costs. In the standard TDF the portfolio allocation rule is a linear function of age only. Under the TTDF the portfolio allocation also depends on the variance risk premium (VRP), which enters as an additional variable in the linear regression. These results are for the case of the investor with risk aversion of 5 (for both funds). The results are reported in percentages.

Maximum Rebalancing	25	25	25	15	15	15
Average Turnover	107	107	107	69	69	69
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	46	38	27	31	24	15
$Std(W_{65})$ (% inc.)	41	32	20	21	13	3
Age-21 Static CE Gain	1.4	1.2	1.0	1.4	1.2	1.0
Age-21 CE Gain	1.4	1.2	1.0	1.4	1.2	1.0
Age-50 Static CE Gain	2.4	2.1	1.7	2.3	2.0	1.6
Age-50 CE Gain	4.4	3.9	3.1	4.0	3.5	2.7
Age-65 Static CE Gain	0.1	−0.1	−0.4	0.1	−0.1	−0.4
Age-65 CE Gain	2.7	2.1	1.2	1.9	1.3	0.5

Table 9: Results with turnover restrictions when $\gamma = 2$

Table 9 presents results from comparing the TTDF with the standard TDF for different rebalancing restrictions and transaction costs. In the standard TDF the portfolio allocation rule is a linear function of age only. Under the TTDF the portfolio allocation also depends on the variance risk premium (VRP), which enters as an additional variable in the linear regression. These results are for the case of the investor with risk aversion of 2 (for both funds). The results are reported in percentages.

Max Rebalancing	25	25	25	15	15	15
Average Turnover	108	108	108	72	72	72
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	56	42	23	44	31	14
$Std(W_{65})$ (% inc.)	35	21	1.2	21	8	−9
Age-21 Static CE Gain	1.4	1.1	0.8	1.4	1.1	0.8
Age-21 CE Gain	1.4	1.1	0.8	1.4	1.1	0.8
Age-50 Static CE Gain	4.6	3.6	2.3	4.5	3.5	2.3
Age-50 CE Gain	13.5	11.0	7.5	12.5	10.1	6.8
Age-65 Static CE Gain	0.1	−0.4	−1.0	0.1	−0.3	−1.0
Age-65 CE Gain	11.4	8.3	4.0	9.1	6.2	2.2

Table 10: Results with less-tight short selling constraints

Table 10 presents results from comparing the TTDF with the standard TDF when both funds are allowed to invest up to 200% in the risky asset. Results are show for different rebalancing restrictions and transaction costs. In the standard TDF the portfolio allocation rule is a linear function of age only. Under the TTDF the portfolio allocation also depends on the variance risk premium (VRP), which enters as an additional variable in the linear regression. These results are for the case of the investor with risk aversion of 5 (for both funds). The results are reported in percentages.

Summary Statistics

Maximum Rebalancing	100	100	100	25	25	25	15	15	15
Average Turnover	376	376	377	116	117	117	73	74	74
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	309	282	242	67	56	40	47	37	24
$Std(W_{65})$ (% inc.)	403	382	350	91	76	56	58	45	28
Age-21 Static CE Gain	2.5	2.3	1.9	2.5	2.3	1.9	2.5	2.3	1.9
Age-21 CE Gain	2.5	2.3	1.9	2.5	2.3	1.9	2.5	2.3	1.9
Age-50 Static CE Gain	4.2	3.7	3.1	3.7	3.3	2.7	3.6	3.2	2.7
Age-50 CE Gain	10.8	9.8	8.3	6.1	5.4	4.4	5.6	4.9	3.9
Age-65 Static CE Gain	0.2	−0.2	−0.7	0.1	−0.1	−0.4	0.1	−0.1	−0.4
Age-65 CE Gain	11.0	9.8	8.1	3.3	2.5	1.5	2.4	1.7	0.8

Table 11: Results with turnover restrictions when $\gamma = 5$

Table 11 presents summary statistics comparing results between the VRP model and the i.i.d. model for the baseline model for different rebalancing restrictions and transaction costs. In the i.i.d. model the investor follows the TDF that conditions only on age and not on the factor, both during working life and during retirement. In the TDFR model the investor follows the TDF that conditions on age and the factor (age coefficients are the same for each factor realization), both during working life and during retirement. Percentage changes reported.

Maximum Rebalancing	25	25	25	15	15	15
Average Turnover	107	107	107	69	69	69
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	58	50.9	40.7	43	36.7	27.7
$Std(W_{65})$ (% inc.)	44	36	24	24	17	8
Age-65 Static CE Gain	2.5	2.2	1.9	2.3	2.1	1.7
Age-65 CE Gain	5.9	5.2	4.3	4.9	4.3	3.4

References

[1] Angerer, X., and P. Lam. “Income Risk and Portfolio Choice: An Empirical Study.” *Journal of Finance*, LXIV (2009), 1037–1055.

[2] Balduzzi, P., and A. W. Lynch. “Transaction Costs and Predictability: Some Utility Cost Calculations.” *Journal of Financial Economics,* 52 (1999), 47–78.

[3] Barberis, N. “Investing for the Long Run when Returns are Predictable.” *Journal of Finance*, 55 (2000), 225–264.

[4] Benzoni, L.; P. Collin-Dufresne; and R. Goldstein. “Portfolio choice over the life-cycle when the stock and labor markets are cointegrated.” *Journal of Finance*, 62 (2007), 2123–2167.

[5] Bollerslev, T., G. Tauchen and H. Zhou, “Expected Stock Returns and Variance Risk Premia,” *Review of Financial Studies*, 22 (2009), 4463–4492.

[6] Bollerslev, T., J. Marrone, L. Xu, and H. Zhou, “Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence, *Journal of Financial and Quantitative Analysis*, 49-3 (2014), 633-661.

[7] Bonaparte, Y.; G. Korniotis; and A. Kumar. “Income Hedging and Portfolio Decisions.” *Journal of Financial Economics*, 113 (2014), 300–324.

[8] Brandt, M. “Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach.” *Journal of Finance*, 54 (1999), 1609–1646.

[9] Brandt, M.W.; A. Goyal; P. Santa-Clara; and J.R. Stroud. "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning about Return Predictability." *Review of Financial Studies*, 18 (2005), 831–873.

[10] Brennan, M.; E. Schwartz; and R. Lagnado. "Strategic Asset Allocation." *Journal of Economic Dynamics and Control*, 21 (1997), 1377–1403.

[11] Brennan, M., and Y. Xia. "Dynamic Asset Allocation under Inflation." *Journal of Finance,* 57 (2002), 1201–38.

[12] Buraschi, A., P. Porchia, and F. Trojani, "Correlation risk and optimal portfolio choice," *The Journal of Finance* 2010 (65), 393–420.

[13] Campbell, J. Y., and L. Viceira. "Consumption and Portfolio Decisions When Expected Returns are Time Varying." *Quarterly Journal of Economics*, 114 (1999), 433–495.

[14] Campbell, J. Y., and L. Viceira. Strategic Asset Allocation. 2002. Oxford University Press.

[15] Campbell, J. Y.; J. Cocco; F. Gomes; P. Maenhout; and L. Viceira. "Stock Market Mean Reversion and the Optimal Allocation of a Long Lived Investor." *European Finance Review*, 5 (2001), 269–292.

[16] Campbell, J.; Y. L. Chan; and L. Viceira. "A Multivariate Model of Strategic Asset Allocation." *Journal of Financial Economics*, 67 (2003), 41–80.

[17] Carroll, C. "Buffer Stock Saving and the Life Cycle / Permanent Income Hypothesis." *Quarterly Journal of Economics,* CXII (1997), 3–55.

[18] Chacko, G., and L. Viceira, "Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets," *Review of Financial Studies*, 18 (2005), 1369–1402.

[19] Cocco, J.; F. Gomes; and P. Maenhout. "Consumption and Portfolio Choice over the Life–Cycle." *The Review of Financial Studies*, 18 (2005), 491–533.

[20] Cooper, R., and G. Zhu. "Household Finance over the Life-Cycle: What does Education Contribute?" *Review of Economic Dynamics*, 20 (2016), 63-89.

[21] Dahlquist, Magnus, Offer Setty and Roine Vestman. "On the Asset Allocation of a Default Pension Fund." *Journal of Finance*, forthcoming.

[22] Davis, S.; F. Kubler; and P. Willen. "Borrowing Costs and the Demand for Equity over the Life Cycle." *The Review of Economics and Statistics,* 88 (2006), 348–362.

[23] Deaton, A. "Saving and Liquidity Constraints." *Econometrica,* 59 (1991), 1221–1248.

[24] Donaldson, S.; F. Kinniry; R. Aliaga-Diaz; A. Patterson; and M. DiJoseph. "Vanguard's Approach to Target-Date Funds." (2013).

[25] Epstein, L., and S. Zin. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57 (1989), 937–969.

[26] Fagereng, Andreas, Charles Gottlieb and Luigi Guiso, "Asset Market Participation and Portfolio Choice over the Life-Cycle," *Journal of Finance,* 72-2 (2017), 705-750.

[27] Fleming, J., C. Kirby, and B. Ostdiek, "The economic value of volatility timing," *Journal of Finance*, 56 (2001), 329–352.

[28] Fleming, J., C. Kirby, and B. Ostdiek, "The economic value of volatility timing using 'realized' volatility", *Journal of Financial Economics* , 67 (2003), 473–509.

[29] Floden, M. "A Note on the Accuracy of Markov-chain approximations to highly persistent AR(1) processes." *Economics Letters*, 99 (2008), 516–520.

[30] Gomes, F., and A. Michaelides, "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *Journal of Finance*, 60 (2005), 869–904.

[31] — and V. Polkovnichenko. "Optimal Savings with Taxable and Tax-Deferred Accounts." *Review of Economic Dynamics*, 12 (2009), 718–735.

[32] Guiso, L. Sapienza, P. and L. Zingales, "Trusting the Stock Market," *Journal of Finance* 63 (2008), 2557-2600.

[33] Haliassos, M., and A. Michaelides. "Portfolio Choice and Liquidity Constraints." *International Economic Review*, 44 (2003), 144–177.

[34] Heaton, J., and D. Lucas. "Portfolio Choice in the Presence of Background Risk." *The Economic Journal*, 110 (2000), 1–26.

[35] Johannes, M„ A. Korteweg, and N. Polson, "Sequential learning, predictability, and optimal portfolio returns," *The Journal of Finance* 69 (2014), 611–644.

[36] Kim, T. S., and E. Omberg. "Dynamic Nonmyopic Portfolio Behavior." *Review of Financial Studies*, 9 (1996), 141–161.

[37] Koijen, R.; T. Nijman; and B. Werker. “When Can Life Cycle Investors Benefit from Time-Varying Bond Risk Premia?” *The Review of Financial Studies*, 23 (2010), 741–780.

[38] Lan, C. “An Out-of-Sample Evaluation of Dynamic Portfolio Strategies.” *Review of Finance*, 19 (2015), 2359–2399.

[39] Lettau, Martin, and Stijn Van Nieuwerburgh, 2008, “Reconciling the return predictability evidence,” *Review of Financial Studies*, 21 (2008), 1607–1652.

[40] Liu, J., “Portfolio selection in stochastic environments,” *Review of Financial Studies*, 20 (2007), 1–39.

[41] Lusardi, A. and O. Mitchell. “The Economic Importance of Financial Literacy: Theory and Evidence.” *Journal of Economic Literature*, 52 (2014), 5-44.

[42] Lynch, A., and S. Tan. “Labor Income Dynamics at Business-Cycle Frequencies: Implications for Portfolio Choice.” *Journal of Financial Economics*, 101 (2011), 333–359.

[43] Michaelides, A. and Y. Zhang, “Stock Market Mean Reversion and Portfolio Choice over the Life Cycle,” *Journal of Financial and Quantitative Analysis*, 52-3 (2017), 1183-1209.

[44] Moreira, A., and T. Muir, “Volatility-managed portfolios,” *Journal of Finance* 72 (2017), 1611-1644.

[45] Moreira, A., and T. Muir, “Should Long-Term Investors Time Volatility?” *Working paper* 2017.

[46] Munk, C., and C. Sorensen. "Dynamic asset allocation with stochastic income and interest rates." *Journal of Financial Economics*, 96 (2010), 433–462.

[47] Pastor, L., and R. F. Stambaugh. "Are Stocks Really Less Volatile in the Long Run?" *The Journal of Finance*, 67 (2012), 431–477.

[48] Pettenuzzo, D.; A. Timmermann; and R. Valkanov. "Forecasting Stock Returns Under Economic Constraints." *Journal of Financial Economics*, 114 (2014), 517–553.

[49] Tauchen, G., and R. Hussey. "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models." *Econometrica*, 59 (1991), 371–396.

[50] Sialms, C., L. Starks and H. Zhang, "Defined Contribution Pension Plans: Sticky or Discerning Money?", Working paper, U. Texas at Austin 2013.

[51] Viceira, L. "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income." *The Journal of Finance*, 55 (2001), 1163–1198.

[52] Vissing-Jorgensen, A. "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution." *Journal of Political Economy*, 110 (2002), 825–853.

[53] Wachter, J. "Optimal Consumption and Portfolio Allocation with Mean-Reverting returns: An exact solution for complete markets." *Journal of Financial and Quantitative Analysis*, 37 (2002), 63–91.

[54] Weil, P. "Non-Expected Utility in Macroeconomics." *Quarterly Journal of Economics*, 2 (1990), 29–42.

[55] Xia, Y. "Learning about Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation." *Journal of Finance*, 56 (2001), 205–246.

Figure 1: Implied volatility (IV), realized volatility (RV) constructed from daily US CRSP returns stock market data and the variance risk premium (VRP) as the difference between the two series. All data are quarterly between 1990 and 2016.

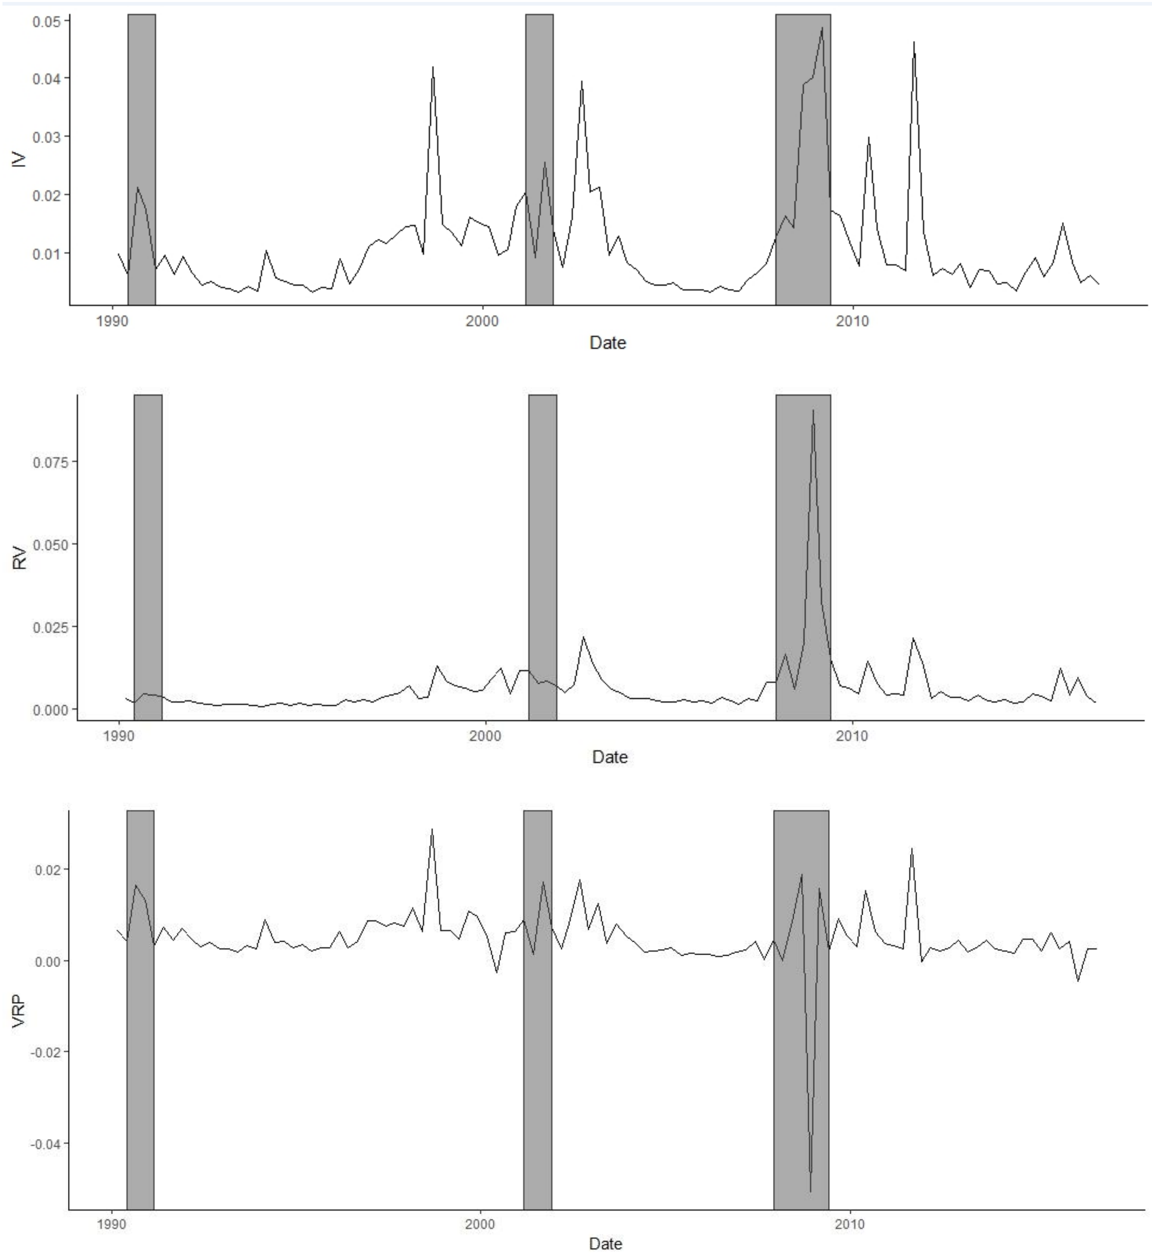

Figure 2 shows over the working part of the life cycle the share of wealth in stocks when the factor is at its median factor realization (factor = 0.49%) in the VRP model, the mean share of wealth in stocks in the VRP model and the mean share of wealth in stocks in the i.i.d. model. The baseline preferences are Epstein-Zin with a risk aversion of 5 and an elasticity of substitution equal to 0.5 and a quarterly discount factor equal to 0.99. VRP is the variance risk premium model and the decision frequency is quarterly.

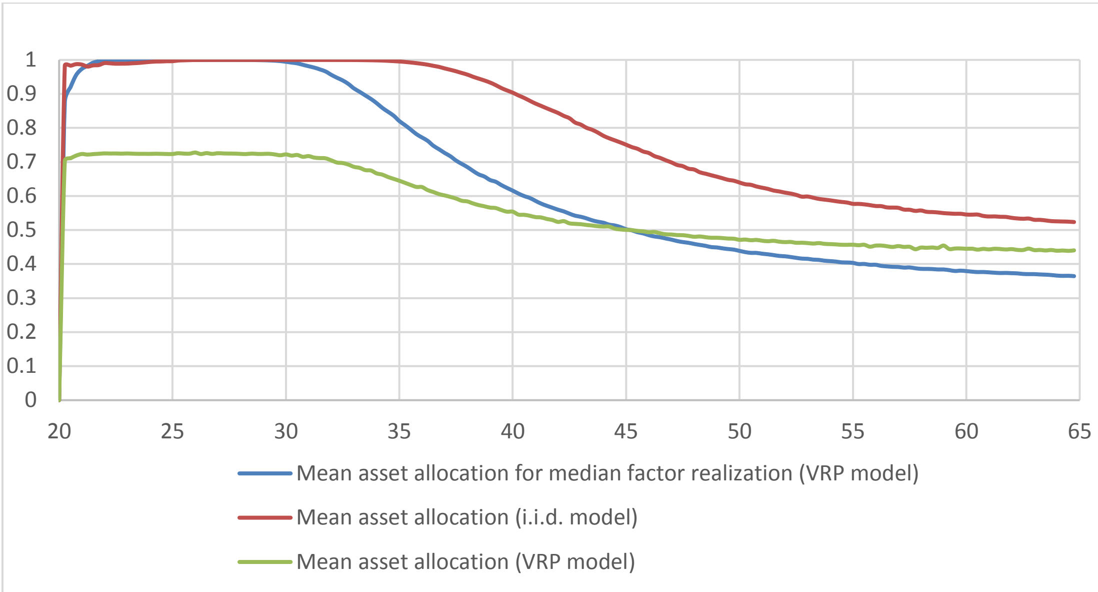

Figure 3 shows the expected portfolio return between the VRP model and i.i.d. model and their difference. The baseline preferences are Epstein-Zin with a risk aversion of 5 and an elasticity of substitution equal to 0.5 and a quarterly discount factor equal to 0.99. VRP is the variance risk premium model and the decision frequency is quarterly.

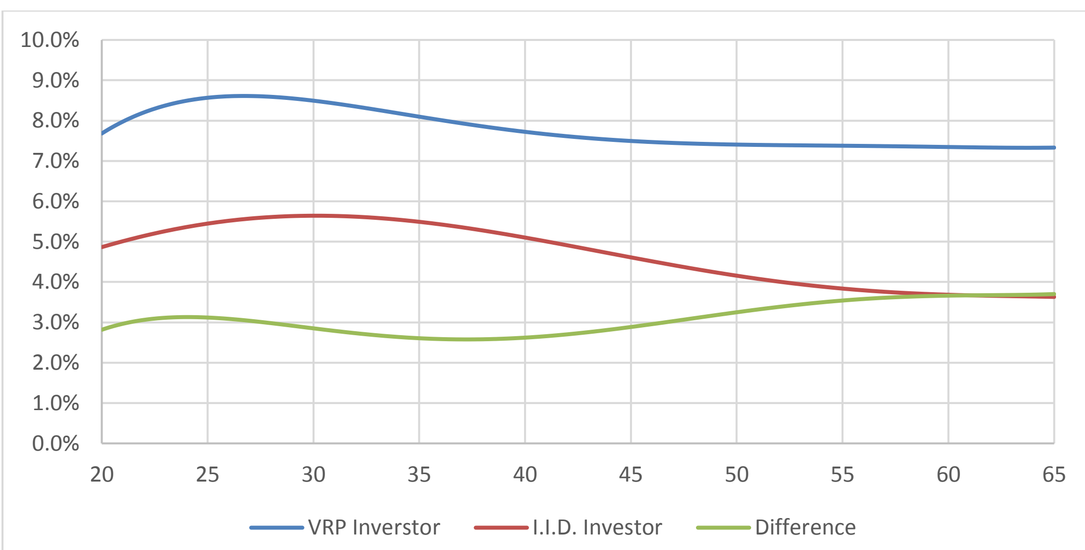

Figure 4 shows the wealth accumulation over the working life (converted to an annual frequency) for the baseline VRP model. The baseline preferences are Epstein-Zin with a risk aversion of 5 and an elasticity of substitution equal to 0.5 and a quarterly discount factor equal to 0.99. VRP is the variance risk premium model and the decision frequency is quarterly.

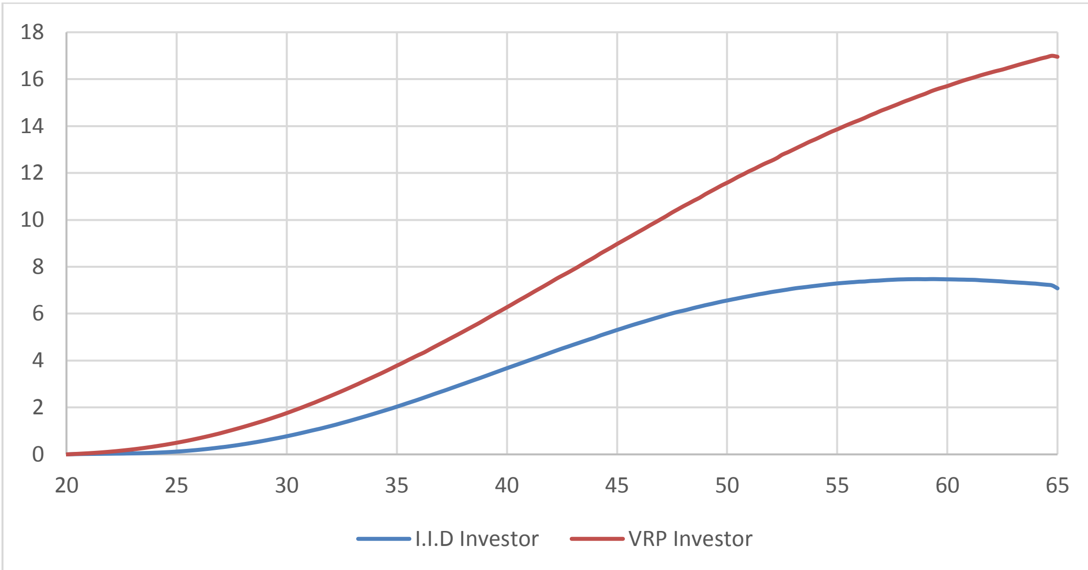

Figure 5 shows the mean share of wealth in stocks for the VRP and i.i.d. models and the target date funds (TDFs) that are constructed based on simulated shares of wealth in stocks and a multivariate regression on age and factor. In the i.i.d. model the factor state is irrelevant (as it should be). The data generating process (DGP) for stock returns in the simulation is the VRP for either model. The baseline preferences are Epstein-Zin with a risk aversion of 5 and an elasticity of substitution equal to 0.5 and a quarterly discount factor equal to 0.99. VRP is the variance risk premium model and the decision frequency is quarterly.

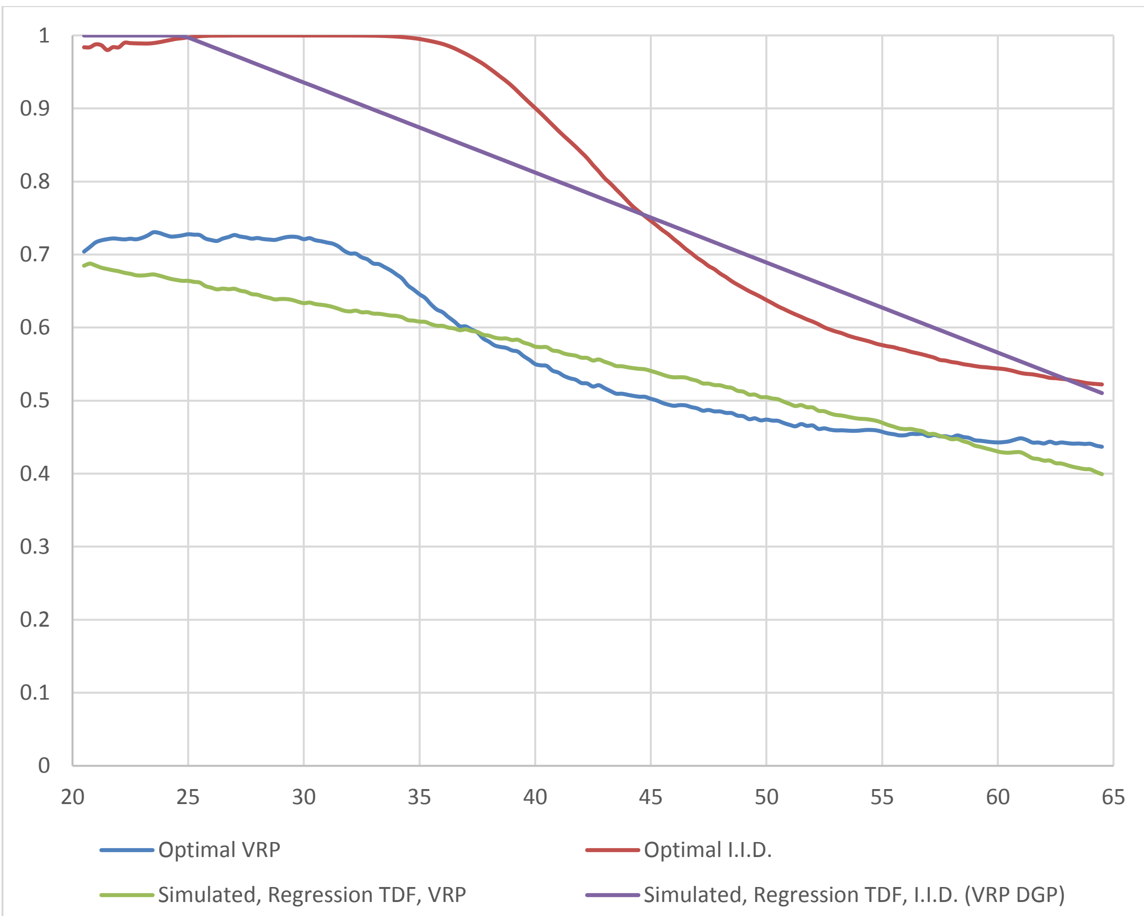

Figure 6 shows the share of wealth in stocks for the target date funds (TDFs) based on different factor realizations and the mean share of wealth in stocks for the VRP model. The data generating process (DGP) for stock returns in the simulation generating the simulated shares of wealth on which the TDF regressions are based is the VRP baseline model. The baseline preferences are Epstein-Zin with a risk aversion of 5 and an elasticity of substitution equal to 0.5 and a quarterly discount factor equal to 0.99. VRP is the variance risk premium model and the decision frequency is quarterly.