Tactical Target Date Funds[*]

Francisco Gomes[†]   Alexander Michaelides[‡]   Yuxin Zhang[§]

June 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

We show that saving for retirement in target date funds (TDFs) modified to take advantage of predictability in excess returns driven by the variance risk premium generates economically large welfare gains. We call these funds tactical target date funds (TTDFs). To be easily implementable and communicated to investors, the portfolio rule followed by TTDFs is designed to be extremely simplified relative to the optimal policy rules. Despite this significant mis-specification, substantial welfare gains persist. Importantly, these gains remain economically important even after we introduce restrictions that limit turnover to empirically observed magnitudes for mutual funds, and after we take into account potential increases in transaction costs. Crucially, we show that this predictability is not correlated with individual household risk, confirming that households are in a prime position to exploit this premium.

JEL Classification: G11, D14, D15

Key Words: Target date funds, life cycle portfolio choice, retirement savings, 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 advice is given by several financial planning consultants (for instance, Vanguard[1]) 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 retirement) 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)).[2] 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 whether exploiting time variation in expected returns can significantly enhance the strategic asset allocation perspective of a life cycle investor saving for retirement, through tactical asset allocation movements over a quarterly frequency.[3,4] 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

---

[1]See Donaldson, Kinniry, Aliaga-Diaz, Patterson and DiJoseph (2013).

[2]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.

[3]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.

[4]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 Moreira and Muir (2017a and 2017b) consider time variation in volatility, while Buraschi, Porchia and Trojani (2010) incorporate time-varying correlations.

can be 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 the current 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 (Lettau and Ludvigson (2001)) or the dividend-yield, capture lower frequency movements (both are more persistent than the VRP) and tend to be associated with bad economic conditions and/or discount rate shocks, both of which might affect households directly.[5] On the other hand, the VRP predictability 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 unlikely to be significantly correlated with household-level risks.

We make this argument empirically by presenting evidence from the Consumer Expenditure Survey (CEX). Specifically, we document that states of the world with high realizations of the VRP do not predict decreases in the household consumption growth, either in the near or in the distant future. Furthermore, they do not predict increases in cross-sectional consumption risk as captured by the cross-sectional standard deviation, skeweness or kurtosis. Moreover, this conclusion holds regardless of whether we condition on stockholder or non-stockholder status, thereby showing that the results do not arise from not conditioning on the stock market participation status (Vissing-Jorgensen (2002)). Importantly, we also show that this holds even though stockholders are shown to bear a disproportionate amount of long run consumption risk as in Malloy, Moskowitz and Vissing-Jorgensen (2009).

As a result of this evidence we can conclude that households are in a prime position to "take the other side" and exploit this premium. Furthermore, in general equilibrium, as households own the financial intermediaries, this adds a further motivation to take the other side of this trade. If those institutional investors are forced to scale down their risky positions when VRP is high because of exogenous constraints, then households should be

---

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

keen to offset this by increasing the risk exposure in their individual portfolios. Perhaps equivalently, a deterministic life style fund completely ignores market information that a sophisticated household might choose to use.

One contribution of our paper is to show how a life cycle model at a quarterly frequency generates similar quantitative insights to the more traditional annual frequency models solved in the literature. Nevertheless, we do not focus on quantifying the welfare gains from following an optimal policy.[6] Instead, we use the output of the 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 individual investors are increasingly expected to be the ones to decide where to allocate their retirement savings, and several of them have limited financial literacy and might be skeptical about complex financial products.[7] Furthermore, we show that despite being an approximate portfolio rule, the rule is still able to capture a significant fraction of welfare gains implied by the optimal policy functions from the model.

We start our argument by showing that relative to an investor who assumes i.i.d. expected returns, the investor exploiting VRP predictability (the VRP investor) earns a significantly higher expected return. This result holds even in the presence of fully binding short-selling constraints which limit the ability of the VRP investor to exploit the time variation in the risk premium. Her expected return in such a model is still between 2.5% to 4% higher at each age (annually). Having identified the optimal portfolio rules and large implied difference in expected returns within the model, we turn to the main question of our paper. Can we design 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

---

[6]Michaelides and Zhang (2017) incorporate stock market predictability through the dividend-yield and compute the welfare gains in the context of a life-cycle model of consumption and portfolio choice.

[7]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.

levels of wealth (relative to future labor income).[8] Furthermore, 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 (TTDF). We estimate the best linear rule from regressions on our simulated data, where we include as explanatory factors not only age, but also the predictive factor (i.e. the variance risk premium).[9] We further truncate the fitted linear rule by imposing fully binding short-sale constraints. We do this because 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 certainty equivalent welfare gains. In our analysis we take into account a potential increase in transaction costs implied by the additional trading implied by the VRP strategy. Even with a quarterly 0.25% decrease in expected returns due to increased portfolio turnover, the certainty equivalent gain from the TTDF versus the standard TDF is still 26% for our baseline calibration. The expected age-65 wealth accumulation is 131% higher. Consistent with the previous results, we find that the gains are particular higher for investors with moderate or high risk aversion, essentially households with higher saving. From this analysis we can conclude that if the TTDFs are introduced, then these investors 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 where we explicitly restrict quarterly turnover to a maximum threshold. It is particularly interesting to discuss the case where we set this threshold so that the average turnover of the constrained TTDF is comparable to the average turnover of the typical mutual fund (78% from Sialms, Starks and Zhang (2013)). Although the increases

---

[8]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.

[9]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.

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, although smaller, remain economically meaningful. For the baseline parameter values the certainty equivalent gain from the TTDF is still 4%.

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 as well. Despite the improved results we believe that all of the above face non-trivial implementation problems relative to the simpler TTDF, and therefore we only present them as extensions to our baseline case.

The paper is organized as follows. Section II discusses the VRP measurement and the VAR model for stock returns. In Section III we show that high realizations of the VRP are not associated with increased household risk. Section IV outlines the life-cycle model and discusses the optimal policy functions. Section V discusses the design of the proposed TTDFs and Section VI explores different extensions. Finally, section VII provides concluding remarks.

2   Variance Risk Premium and Stock Returns

2.1   VAR model for stock returns

The time variation in expected returns is captured by a predictive factor (f_t) and following 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. Following Bollerslev, Tauchen and Zhou (2009) and Boller-slev, Marrone, Xu, and Zhou (2014) we consider the variance risk premium (VRP) as the predictive factor, i.e. $f_t \equiv VRP_t$. The formulation allows for contemporaneous correlations between z_{t+1} and ε_{t+1}.[10]

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

$$r_{t+1} - r_f = \mu + z_{t+1}. \tag{3}$$

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 predictability. This will be implemented by adjusting appropriately the value of α in equation (1).

2.2 Variance Risk Premium

As in Bollerslev, Tauchen and Zhou (2009) we define the variance risk premium (VRP_t) as the difference between the option-implied variance of the stock market (IV_t) and its realized variance (RV_t),

$$VRP_t \equiv IV_t - RV_t. \tag{4}$$

The data for the quarterly implied variance index (IV_t) are taken from the Federal Reserve Bank of St. Louis (FRED) while the data for the monthly realized variance (RV_t) from Zhou (2017).[11] We convert the monthly realized variance to quarterly by adding the monthly terms. Figure 1 shows the time series variation in implied variance (IV_t), realized variance (RV_t) and the variance risk premium (VRP_t). Figure 1 replicates and extends essentially the original Bollerslev, Tauchen and Zhou (2009) measure.

---

[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]Available here https://sites.google.com/site/haozhouspersonalhomepage/.

2.3   VAR Estimation

Table 1 contains the descriptive statistics from the data set. 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. The net constant real interest rate, r_f, equals 0.37% corresponding to 1.5% at an annual frequency.

Table 2 reports the estimation results for the VAR model (equations (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 the unconditional variance of unexpected stock market returns from

$$\sigma_z^2 = Var(r_t) - \beta^2 \sigma_f^2 \tag{5}$$

The correlation between the factor and the return innovation ($\rho_{z,\varepsilon}$) is a potentially important parameter in determining 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.[12]

3   VRP and Household Consumption Risk

3.1   Discussion

The empirical results in the previous section document that a high value of the VRP forecasts high expected stock returns next quarter, consistent with the findings in Bollerslev et al. (2009). However, the optimality of increasing the allocation to stocks when the VRP is high will be over-stated if the high expected returns next quarter are accompanied by an increase

---

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

in risk for households. Therefore, it becomes important for our analysis that this is not the case, and in this section we provide the corresponding supporting evidence.

It is important to clarify that we are not arguing that the changes in expected returns forecasted by the VRP do not reflect risk, as such a discussion is beyond the scope of our paper. We are merely stating that, if it is indeed risk, this risk appears to be faced primarily by other agents in the economy and not by individual households directly. For example, institutional investors such as mutual funds or banks face constraints that might lead them to reduce their risk bearing capacity in these periods.[13] If households are not directly exposed to this risk, it is therefore natural for them to increase their allocation to stocks in these periods and thus earn the additional premium by effectively taking the other side of this trade.[14] Furthermore, from a general equilibrium perspective, and to the extent that it is the same households that own the banks and therefore their own wealth that is invested in pension/mutual funds, a further motivation arises for taking the other side of the VRP. As institutional investors are forced to scale down their risky positions, then households should be keen to offset this by increasing the risk in their individual portfolios.

3.2 Consumer Expenditure Survey

We use non-durable consumption and services from the Consumer Expenditure Survey (CEX) following primarily the methodology in Malloy et. al. (2009).[15] We construct quarterly consumption growth rates for stockholders and non-stockholders from January 1996 to December 2015. The CEX is a repeated cross section with households interviewed monthly over five quarters, enabling us to compute quarterly growth rates at a monthly frequency. Nevertheless, we cannot follow the same household for more than five quarters, and therefore membership in a group is used to create a pseudo-panel to track household risk over longer time periods. Following the literature, we regress the change in log consumption on drivers not in the model (log family size and seasonal dummies) and use the residual as our quarterly

---

[13]For example, tracking error constraints for mutual funds or VAR constraints for banks.

[14]Naturally if we take the view that a high value of the VRP does not represent an increase in risk at all, then the same conclusion applies: households should exploit this predictable variation in the risk premium.

[15]An internet appendix provides further details on data construction than the appendix at the end of this paper.

consumption growth measure.

Our model applies primarily to stockholders and we know from prior theoretical and empirical work that stockholders face different risks from non-stockholders. We therefore estimate separate measures of risk for the two groups. To determine stockholders we use the financial information provided in interview five and we also drop any households for which any of the interviews in the second to fifth quarter are missing. To determine stockholder status we use the response to the question of owning "stock, bonds, mutual funds and other such securities".

We first show that the main insight from Malloy et. al. (2009), namely that long run risks matter for stockholders, holds for our updated data set. Specifically, Table A1 in our appendix replicates Table I from their paper that documents the sensitivity of stockholder and nonstockholder consumption growth to aggregate consumption growth from NIPA over different horizons. Specifically, we compute the average consumption growth rate for a particular group of households (for instance, stockholders) for different horizons s=1, 2, 12, 24, by averaging the log consumption growth rates as

$$\frac{1}{N} \sum_{i=1}^{N} [c_{i,t+s} - c_{i,t}] \tag{6}$$

where N can vary depending on group and time period and $c_{i,t}$ is the quarterly log consumption of household i at time t.

Table A1 in the appendix shows that the coefficient from regressing a group's discounted consumption growth over horizon $s = 1, 2, 12, 24$ on aggregate discounted[16] consumption growth over the same horizon generates different conclusions across groups. Specifically, stockholder consumption growth is more sensitive to aggregate consumption growth than nonstockholder consumption growth. Moreover, the differences are even larger at longer horizons, supporting the interpretation that stockholders bear a much larger proportion of stock market risk (Malloy et. al. (2009)). From our perspective, we would like to understand in the next subsection whether exposure to VRP risk can have a similar interpretation.

---

[16]We use a discount factor equal to one to be more comparable with regressions we produce later but the results are quantitatively similar with a discount factor less than one.

3.3 Empirical Evidence

How can we determine whether consumption growth risks in the short run and long run are affected by the VRP? We address this question by reporting results from regressing different moments of cross-sectional consumption growth on the VRP.

We start by considering regressions of mean consumption growth at different horizons and for different groups on the current VRP. More precisely, we perform the following regressions

$$\frac{1}{N} \sum_{i=1}^{N} [c_{i,t+s} - c_{i,t}] = \alpha_c^s + \beta_c^s * VRP_t + \varepsilon_t^c, \ s = 1, 2, 12, 24 \tag{7}$$

As discussed above, given the nature of the CEX, we can only compute consumption growth rates for the same agent for up to $s = 2$. However, motivated by the long-run risk literature and by the evidence in Malloy et. al. (2009), we also consider the possibility that a high variance risk premium might signal an increase in long-run consumption risk by investigating the statistical significance of (β_c^s) in equation (7) as the horizon s increases.

The estimates of β_c^s are shown in Panel A of Table 3. The standard errors are computed using a Newey-West estimator that allows for autocorrelation of up to $s-1$ lags when $s > 1$. For both stockholders and non-stockholders, β_c^1 is non-significant, indicating that a high value of VRP is not associated with lower expected future consumption growth rate in the next quarter. The same conclusion is obtained for $s = 2$ and the same conclusion arises as we consider consumption growth rates over multiple years ($s = 12$ and $s = 24$). We conclude that there is no significant relationship between VRP and individual short run or long run household consumption risk for either stockholders or non-stockholders.

We can repeat the same analysis for higher cross sectional moments of consumption growth rates. Nevertheless, because higher moments (the standard deviation, skewness and kurtosis) are not additive like the mean consumption growth rate, we can only report the regressions for consumption growth rates for $s = 1$ and $s = 2$. This is reported in Panels B, C, and D, of Table 3 where we explore the possibility that the VRP might be associated with a future increase in cross-sectional consumption risk. We find that high VRP states are not associated with an increase in either the cross-sectional standard deviation of consumption

10

growth or its kurtosis, or with a decrease in its skewness. We conclude that, given the lack of any statistical significance in these regressions, high VRP states are not associated with an increase in future cross-sectional household consumption risk.

Overall, our results confirm that high VRP states, while predicting high future expected returns, are on average not followed by periods of lower household consumption growth or high cross-sectional dispersion in consumption growth rates.

4 Life-Cycle Asset Allocation Model

Time is discrete, but contrary to most of the life-cycle asset allocation 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, for a total of 324 quarters. In the notation below we will use t to denote calendar time and a to denote age.

4.1 Preferences and Budget Constraint

In the model there are two financial assets available to the investor. The first one is a riskless asset representing a savings account. 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 potentially time varying as captured by the VAR model described in section 2 (equations (1) and (2)).

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),

$$V_a = \max \left\{ (1-\beta)C_a^{1-1/\psi} + \beta \left(p_a E_a(V_{a+1}^{1-\gamma}) \right)^{\frac{1-1/\psi}{1-\gamma}} \right\}^{\frac{1}{1-1/\psi}}, \tag{8}$$

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

$$W_{a+1} = [\alpha_a R_{t+1} + (1 - \alpha_a)R_f](W_a - C_a) + (1 - h_{a+1})Y_{a+1} \tag{9}$$

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

$$\alpha_a \in [0, 1] \tag{10}$$

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.

4.2 Labor Income Process

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[17]

$$Y_a = \exp(g(a))Y_a^p U_a, \tag{11}$$

$$Y_a^p = Y_{a-1}^p N_a \tag{12}$$

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

---

[17]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.

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).

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 denote the normalized variables the dynamic budget constraint becomes

$$w_{a+1} = \frac{1}{N_{a+1}}[r_{t+1}\alpha_{ia} + r_f(1 - \alpha_{ia})](w_a - c_a) + (1 - h_{a+1})\exp(g(a+1))U_{ia+1}. \qquad (13)$$

As 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

$$Y_a = \lambda Y_K^p \text{ for } a > K \qquad (14)$$

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

4.3   Estimation and Calibration

We take the deterministic component of labor income ($g(a)$) from the estimates in Cocco et al. (2005) 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.

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 as there is no available empirical guidance on this parameter.

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.

We use preference parameters previously used (Gomes and Michaelides (2005)) or estimated (Cooper and Zhu (2016)) in the literature using U.S. data. The discount factor is 0.9875 (annual equivalent around five percent), the elasticity of intertemporal substitution equal to 0.5 and relative risk aversion coefficient equal to 5.0. We have undertaken extensive comparative statics around these parameters that we do not report to keep the paper concise.

4.4 Optimal portfolio allocation

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). These results will form the basis for the next section, where we propose the tactical target date funds (TTDFs). 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)).[18] 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

$$\alpha_a[E(f)] < E[\alpha_a^{iid}] \tag{15}$$

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.[19]

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

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

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, the optimal allocation of the VRP investor is sometimes far below $\alpha_a[E(f)]$ and never exceeds it by much.[20]

---

[18]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.

[19]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.

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

Building on the previous intuition, it is not surprising to find that the sign of inequality flips once the portfolio allocation at the mean factor realization ($\alpha_a[E(f)]$) falls below 50%, which takes place around age 45. Now the more binding constraint is the short-selling constraint on stocks so we have:

$$E[\alpha_a(f)] > \alpha_a[E(f)] \text{ for } a > 45 \tag{17}$$

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 inequalities (15) and (16) it is easy to see that, until age 45, we have:

$$E[\alpha_a(f)] < E[\alpha_a^{iid}] \tag{18}$$

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 (16) is replaced by (17), we see that, although the difference between the optimal allocation of the VRP and i.i.d. investors decreases, equation (18) still holds: inequality (15) dominates inequality (16).

4.5   Portfolio returns

In this section we study the differences in expected returns between the VRP and i.i.d. investors. To avoid repetition we ignore transaction costs in these calculations, since 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 at each age

$$E(R_{t+1}^P) = \alpha_a E_t[R_{t+1}] + (1 - \alpha_a)R_f, \, a = 1, ..., T \tag{19}$$

which are computed by averaging (at each age) across all simulations.

16

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.

5   Tactical Target-date Funds

In the previous section we derived the optimal life-cycle policy functions from the model. However, these are not feasible options for a mutual fund. For example, current 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, using a roughly linear or piece-wise 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, even though 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 has the dual advantage of being simple to explain and a reasonable approximation to an heterogeneous set of optimal life-cycle profiles. 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 decide where to allocate their retirement savings.

In the same spirit, and in our 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 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. On the other hand, the more complicated rules are more likely to suffer from over-fitting or model misspecification. 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 Designing Tactical 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

$$\alpha_{iat} = \theta_0 + \theta_1 * a + \theta_2 * f_t + \varepsilon_{iat}. \tag{20}$$

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 and Figure 4 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.[21] Table 4 also 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 (20)) 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 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).

---

[21]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.

5.1.2 Tactical TDF conditioning on the VRP (TTDF2)

As previously discussed, the portfolio rule based on equation (20) 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 regressions conditional on the different realizations of the predictive factor. That is, we run the following series of regressions for each f_j in our discretization grid

$$\alpha_{iat} = I_{f_t=f_j} \theta_0^j + I_{f_t=f_j} * \theta_1^j * a + \varepsilon_{iat}^j, \text{ for each } f_j \tag{21}$$

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 5. Table 5, Panel A reports for the baseline case of risk aversion equal to 5 the regression results for three different values of f_j : the mean, plus two and minus two standard deviations of the factor (VRP).[22] Panels B and C report the same results for risk aversions of 2 and 10, respectively. As we can see, a 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 (20)) 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 from Tactical Target Date Funds

5.2.1 Welfare Metric

Having identified a feasible portfolio rule for the TTDF we now proceed to compute the corresponding certainty-equivalent utility gains. Consistent with the focus of our paper to design improved target date funds, the baseline welfare calculations are computed by keeping pre-retirement consumption constant and comparing age-65 certainty equivalents, following Dahlquist, Setty and Vestman (forthcoming). The differences in certainty equivalents therefore represent the increase or decrease in *risk-adjusted* consumption levels that

---

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

the agent will register during the retirement period. This procedure guarantees that the pre-retirement utility is the same across cases (TDF and TTDF) and therefore the certainty equivalent gain at retirement captures the full welfare change.

In comparing different rules we assume the same asset allocation rules after 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 the TDF only (that is, during working life). 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 is able to identify the fund that matches her level of risk aversion, both for the TTDFs and the standard TDF.

5.2.2 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 (equation (21)) 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. In these calculations, as previously mentioned, we also take into account a potential increase in transaction costs implied by the market timing strategy. More specifically, we take into account 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 is, the target fund 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 (only) by the TTDF2.[23] The results are reported in Table 6.

We first consider the case with no transaction costs ($tc = 0$). For all three values of risk aversion the increases in wealth accumulation at age 65 are extremely high: 201%,

---

[23]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.

260% and 337%. Likewise, the corresponding age-65 certainty equivalent gains are also very large: 37.8%, 55.3% and 95%, respectively. As we introduce differential transaction costs for the TTDF2 these values naturally fall. However, even for an arguably large quarterly transaction cost of 25 basis points ($tc = 0.25\%$), age-65 wealth is higher by more than 100% for all investors. As a result, the utility gains remain quite high: 38.6% for the baseline risk aversion of 5, increasing (decreasing) to 78.9% (23.5%) for the risk aversion of 2 (10). 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 with high financial literacy who would potentially be willing to invest in such funds if they were introduced, could obtain very large CE gains from doing so.

5.2.3 Tactical Target Date Fund (TTDF)

We now study the results for the simpler TTDF rule (equation (20)). These are shown in Table 7, again for different values of risk aversion (γ) and different values of the additional transaction costs (tc). When considering the case with $tc = 0.0$ the increases in age-65 wealth accumulation are 103%, 182% and 312%, for risk aversion of 2, 5 and 10, respectively. The associated CE gains are 20.3%, 40.5% and 80.3% showing that the simple rule proposed by equation (20) is able to capture extremely large gains. This is particularly remarkable if we recall that, in this analysis, we are assuming that the investor does not exploit the predictability in expected returns at retirement.

Importantly, the welfare gains remain economically large even as we introduce the additional transaction costs. For the baseline calibration of risk aversion (5), even with a 25 basis points increase in costs, relative to those of the standard TDF, age-65 wealth accumulation is still 131% higher under the TTDF and the certainty equivalent consumption gain is 26.2%. As before, these values are even higher for the less risk-tolerant investor (64.4%) and lower for the more risk-tolerant one (10.1%). One implication of these results is that it would be particularly beneficial to introduce the TTDFs in pension plans with investors with moderate or high risk aversion (5 or 10). The important point is that households that have the tendency to be net savers will benefit more from such funds than households with

22

lower saving rates. Equivalently, if such funds are offered in parallel with standard target date funds, investors that save more are the ones that would benefit the most from switching away from the conventional product.

5.3 Introducing Turnover Restrictions

5.3.1 Approach

One potential concern with the TTDFs, as presented in the previous section, is that their 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 investor 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 fund's turnover. The restriction limits the optimal rebalancing of the portfolio share to a maximum threshold (k). More precisely, the portfolio rule is subject to the additional constraint

$$\alpha_a = \begin{cases} \alpha_{a-1} + k & if\ \alpha_a^* > \alpha_{a-1} + k \\ \alpha_a^* & if\ |\alpha_a^* - \alpha_{a-1}| < k \\ \alpha_{a-1} - k & if\ \alpha_a^* < \alpha_{a-1} - k \end{cases} \tag{22}$$

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

In our analysis we consider two thresholds, $k = 25\%$ and $k = 15\%$. We impose equation (22) ex-post on the previously estimated policy rules, instead of solving the corresponding dynamic programming problem for two reasons.[24] First, even though the optimal policy function would by definition satisfy constraint (22), that does not guarantee that the cor-

---

[24]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.

responding fitted linear rule estimated from the simulated data would as well.[25] Second, from an implementation perspective this again makes the rule more transparent and easy to follow and explain to an investor. The asset allocation of the fund is given by the previous regression specification, which yields α_a^*, subject to this constraint.

5.3.2 Results

In Table 8 we show the results when we impose constraint (22), for the baseline 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). High fund turnover was the motivation for including the additional transaction costs in the previous subsection. Therefore, since we are now limiting fund turnover directly, in these results we only consider the cases with $tc = 0.0$ and $tc = 0.10\%$.

The constraints naturally 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 for the baseline case (risk aversion of 5) was 269% in the absence of the turnover constraints, but falls to 45% and 14% for $k = 25\%$ and $k = 15\%$, respectively. However, this is accompanied by an equally significant reduction in the impact on the standard deviation of (age-65) wealth. In the absence of turnover constraints this standard deviation had increased by 462%, whereas now the percentage change is limited to 48% and 5%, respectively.

As we introduce these additional restrictions the extremely large welfare gains that we previously documented are reduced, but we still obtain values that are economically quite meaningful and, we would argue, much more reasonable. With $tc = 0.0$ the certainty equivalent gains for the baseline case (risk aversion of 5) are 11.1% and 3.7%, for $k = 25\%$ and $k = 15\%$, respectively. Even with $tc = 0.1\%$ both of these still remain positive: 7.2% and 0.4%, respectively.

---

[25]This is the same issue we already had before with the short-selling constraints and these also had to be imposed ex-post.

As we consider investors with either higher or lower risk aversion we again find that the certainty equivalent gains are particularly larger for the former. Even with the tighter turnover restriction ($k = 15\%$ and $tc = 0.10\%$), the investor with risk aversion of 10 still accumulates 83% more wealth at age 65, on average, by using the TTDF. This corresponds to a certainty equivalent gain of 22%. Across all cases, the investment in the TTDF only leads to certainty equivalent loss for one them: the combination of the tighter turnover restriction *and* additional transaction costs for the investor with risk aversion 2. But as just discussed, even under this combination the investor with risk aversion of 10 still has a certainty equivalent gain of 22%.

Two of the $tc = 0.0$ cases are particularly interesting: the one for the investor with risk aversion of 2 and $k = 25\%$, and the one for the investor with risk aversion of 5 and $k = 15\%$. In both of these the change in the standard deviation of age-65 wealth is very small, -2% and 5% respectively, yet there are meaningful differences in wealth accumulation: 23% in the first case and 14% in the second. So for a very similar level of ex-ante risk the investor is obtaining a noticeable difference in average expected wealth. This is reflected in certainty equivalent gains of 4.9% and 3.7%.

Overall, the results in Table 8 confirm that it is possible to design a relatively simple target date fund rule that exploits the risk premium predictability obtained from the VRP, while only requiring standard levels of turnover, and being able to generate economically large welfare gains for a wide range of investors, especially the ones that are net savers over the working life cycle.

6 Extensions

6.1 Relaxing the short-selling constraints

As shown in Figure 5, 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. Moreover, a mutual fund that takes leveraged positions might not be regarded as an acceptable choice by some pension plan providers. Nevertheless, the proposed TTDF strategy will be implemented by a mutual fund and hence it should be much cheaper and feasible to take both borrowing and short-selling positions.

In this section we therefore 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:

$$\alpha_a \in [0, 2] \tag{23}$$

For the range of parameter values that we consider the upper bound on this constraint becomes less binding. 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 of 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 9 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.[26]

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 91.8% (40.5%) to 67.3% (26.2%) for $tc = 0$ (0.25%). A less tight short-selling constraint

---

[26]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.

(equation (23)), 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 360% as opposed to 213% for the case with fully binding short-selling constraints. To address this concern Table 9 also reports results with the exogenous constraint on trading (equation (22)). As we introduce the tighter version of constraint ($k = 15\%$) portfolio turnover drops significantly, to around 73%. The welfare gains naturally decrease substantially but, as before, remain economically significant. As we compare them with the ones in Table 8 we find that they are very similar but still higher. For example, for $tc = 0.0$, the certainty equivalents are now 13.2% and 5.2% for $k = 25\%$ and $k = 15\%$, respectively, compared with 11.1% and 3.7% in Table 8.

We conclude that relaxing the short-selling constraint on the riskless asset can increase the welfare gains from investing in the TTDF, even if we restrict the fund's turnover to reasonable levels.

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 run a second regression given by equation (20) for ages greater than 65. From this we obtain a linear portfolio rule for the retirement period which complements the TTDF, that is a TTDF in retirement.

The results are shown in Table 10 for the baseline case of risk aversion 5 and with turnover restrictions to keep trading volume consistent with that of typical mutual funds. As expected, the welfare gains are now even larger. For the tighter turnover restrictions the certainty equivalent gains are between 12.3% and 20.1% substantially larger than the

comparable ones from Table 8 (0.4% to 3.7%, respectively).

7   Conclusion

We analyse how target date funds can combine the long term strategic asset allocation perspective of a life cycle investor with the short term market information that gives rise to tactical asset allocation. We rely on the variance risk premium (VRP) as the main factor producing variation in the expected risk premium in quarterly frequency and embed this in a life cycle model to derive optimal saving and asset allocation. We then show how enhanced funds, which we call Tactical Target Date Funds (TTDFs), can be designed in a parsimonious way and can deliver substantial welfare gains. These gains are substantial and remain economically large even after we include transaction costs and further explicitly restrict the turnover of the TTDF. In unreported experiments we extend the analysis to a wider set of preference parameter configurations and different models of investor behavior during retirement. Further research into the design and commercialization of the proposed TTDFs, and the potential complications that may arise in such implementations, is an interesting topic for future research.

Appendix
Consumer Expenditure Survey

We use non-durable consumption and services aggregated from the CEX and exclude durables, implicitly assuming that utility is separable between durables and non-durables and services. This also allows comparison with earlier literature, particularly Malloy et. al. (2009). The service categories relating to durables are also excluded (housing expenses but not costs of household operations), medical care costs, and education costs as they have substantial durable components.

Our CEX sample choice follows Malloy et. al. (2009). Extreme consumption outliers for which consumption growth is less than 0.2 and greater than 5.0 are dropped. To determine stockholders we use the financial information provided in interview five and we also drop any households for which any of the interviews in the second to fifth quarter are missing. To determine stockholder status we use the response to the category "stock, bonds, mutual funds and other such securities".

Table A1: Sensitivity of Household Consumption Growth to Aggregate Consumption Growth

Table A1 presents the sensitivity of stockholder and non-stockholder consumption growth to aggregate consumption growth taken from NIPA for horizons of S = 1, 2, 12, and 24 quarters. The sensitivity is computed as the regression coefficient from regressing a group's consumption growth over horizon S on current aggregate consumption growth. Below each entry we include the t-stat. Standard errors are computed using a Newey-West estimator that allows for autocorrelation of up to $S - 1$ lags when $S > 1$.

	Mean Consumption Growth			
S	1	2	12	24
Stockholders	1.05	1.12	1.53	0.97
(t-stat)	(0.94)	(1.45)	(4.14)	(4.77)
Non-Stockholders	0.78	0.61	0.94	0.55
(t-stat)	(1.17)	(1.89)	(3.83)	(3.36)

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,* 65, (2010), 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] Chacko, G., and L. Viceira, "Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets," *Review of Financial Studies*, 18 (2005), 1369–1402.

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

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

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

[21] 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.

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

[23] 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.

[24] 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.

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

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

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

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

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

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

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

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

[33] 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.

[34] Lettau, M., and S. Ludvigson, "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56, 3, (2001), 815–849.

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

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

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

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

[39] Malloy, C.J., T.J. Moskowitz, and A. Vissing-Jorgensen, "Long-run stockholder consumption risk and asset returns," *The Journal of Finance*, 64(6), (2009), 2427-2479.

[40] 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.

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

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

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

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

[45] Polkovnichenko, V.; "Life-Cycle Portfolio Choice with Additive Habit Formation Preferences and Uninsurable Labor Income Risk." *The Review of Financial Studies*, 20, 1, (2007), 83-124.

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

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

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

[49] 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.

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

[51] Zhou, H. "Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty." *Annual Review of Financial Economics*, (2017), forthcoming.

Table 1: Descriptive Statistics for Returns and Variance Risk Premium

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. This 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
Mean (%)	1.98	1.11	0.62	0.49	0.60
SD (%)	7.84	0.94	0.98	0.75	0.80
Kurtosis	3.24	8.16	54.23	31.83	9.64
Skewness	-0.40	2.25	6.45	-3.24	-1.39
AR(1)	0.00	0.41	0.47	-0.17	0.001
Panel B: Correlation Matrix					
1990Q1 –2016Q3	r	IV	RV	$IV-RV$	π
r	1.00	-0.52	-0.42	-0.10	-0.11
IV	–	1.00	0.70	0.34	-0.18
RV	–	–	1.00	-0.43	-0.46
$IV-RV$	–	–	–	1.00	0.38
π	–	–	–	–	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: Sensitivity of Household Consumption Growth to VRP Across Horizons

Table 3 presents the sensitivity of different moments of stockholder and non-stockholder consumption growth to the variance risk premium (VRP) over horizons of S = 1, 2, 12, and 24 quarters. Panel A reports the sensitivity of mean consumption growth, while Panels B, C and D report the results for Standard Deviation, Skewness and Kurtosis, respectively. The sensitivity is computed as the regression coefficient from regressing a group's consumption growth over horizon S on current VRP. Below each entry we include the t-stat. Standard errors are computed using a Newey-West estimator that allows for autocorrelation of up to $S-1$ lags when $S > 1$.

Panel A: Mean consumption Growth				
S	1	2	12	24
Stockholders	1.15	0.58	-0.24	0.90
(t-stat)	(1.59)	(0.98)	(0.29)	(0.99)
Non-Stockholders	0.43	0.19	0.50	0.88
(t-stat)	(1.34)	(0.50)	(0.86)	(1.57)

	Panel B: Std. Dev.		Panel C: Skewness		Panel D: Kurtosis	
S	1	2	1	2	1	2
Stockholders	0.39	-0.89	0.96	3.39	-8.37	14.08
(t-stat)	(0.85)	(1.78)	(0.31)	(0.64)	(1.21)	(0.93)
Non-Stockholders	0.18	0.04	2.81	2.27	-4.73	-4.90
(t-stat)	(0.58)	(0.18)	(1.60)	(0.91)	(1.03)	(1.08)

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			$i.i.d.$
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 ($\gamma = 2$)				
	VRP			$i.i.d.$
Factor	2 s.d. above mean	Mean	2 s.d. below mean	
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 ($\gamma = 10$)				
	VRP			$i.i.d.$
Factor	2 s.d. above mean	Mean	2 s.d. below mean	
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			5			10		
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	201	172	134	260	234	198	377	358	331
$Std(W_{65})$ (% inc.)	289	254	205	363	336	297	570	561	546
Age-65 CE Gain	37.8	31.6	23.5	55.3	48.2	38.6	95.0	88.4	78.9

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			5			10		
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	103	83	57	182	161	131	312	293	265
$Std(W_{65})$ (% inc.)	97	77	50	248	223	187	506	491	466
Age-65 CE Gain	20.3	15.9	10.1	40.5	34.4	26.2	80.3	73.7	64.4

Table 8: Results with turnover restrictions

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. The results are reported in percentages.

Risk Aversion	2				5				10			
Max Rebalancing	25	25	15	15	25	25	15	15	25	25	15	15
Mean Turnover	108	108	72	72	106	106	69	69	100	100	66	66
tc (inc.)	0.00	0.10	0.00	0.10	0.00	0.10	0.00	0.10	0.00	0.10	0.00	0.10
W_{65} (% inc.)	23	10	0.6	-11	45	32	14	3.4	141	128	94	83
$Sd(W_{65})$ (% inc.)	-2	-14	-27	-36	48	33	5	-5.5	275	255	190	173
Age-65 CE Gain	4.9	1.8	0.24	-2.5	11.1	7.2	3.7	0.35	38.5	34.0	26.2	22.3

Table 9: Results with less-tight short selling constraints

Table 9 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.

Maximum Rebalancing	100			25			15		
Average Turnover	360	360	360	114	114	116	73	73	74
tc (inc.)	0.00	0.10	0.25	0.00	0.10	0.25	0.00	0.10	0.25
W_{65} (% inc.)	572	523	452	65	44	16	24	7	-15
$Std(W_{65})$ (% inc.)	615	599	569	108	82	49	41	21	-2.5
Age-65 CE Gain	91.8	81.5	67.3	13.5	8.1	1.1	5.2	0.6	-5.0

Table 10: Results for exploiting predictability both during working life and retirement

Table 10 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. The portfolio allocations of both the iid and the VRP investors are given by the corresponding funds both during working life, TDF and TTDF respectively, and during retirement. The asset allocations of the retirement funds are constructed following the same procedure as for the pre-retirement funds. Percentage changes reported.

Maximum Rebalancing	25		15	
Average Turnover	108	108	69	69
tc (inc.)	0.00	0.10	0.00	0.10
W_{65} (% inc.)	63	41	34	13
$Std(W_{65})$ (% inc.)	107	71	55	22
Age-65 CE Gain	30.0	21.6	20.1	12.3

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.
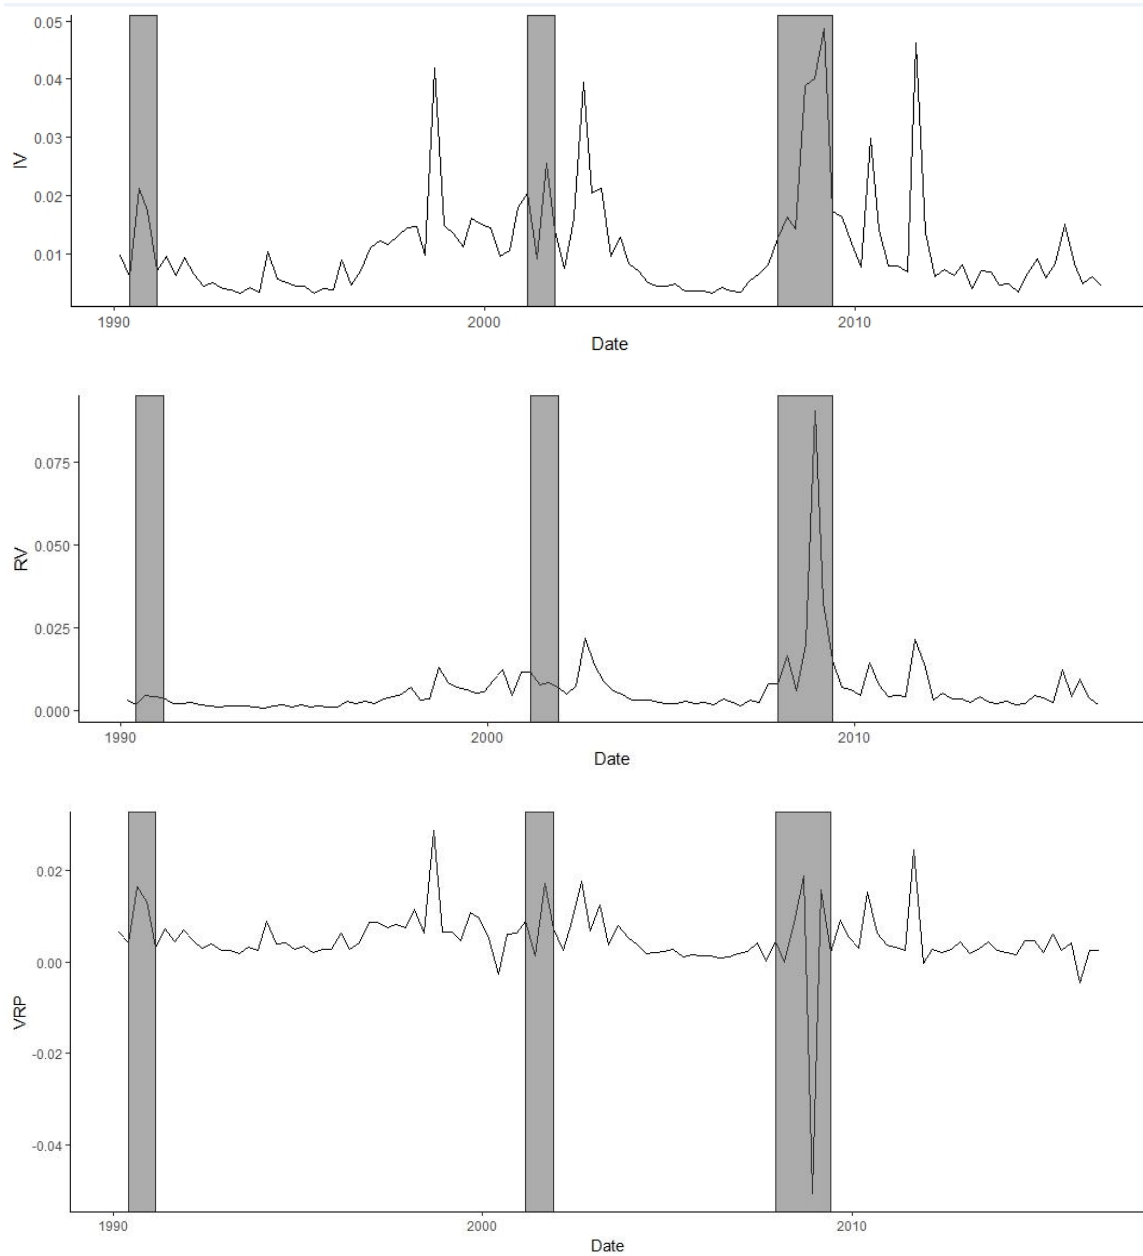
Figure 2 shows over the working part of the life cycle the share of wealth in stocks when the factor is at its median factor realization (factor = 0.49%) in the VRP model, the mean share of wealth in stocks in the VRP model and the mean share of wealth in stocks in the i.i.d. model. The baseline preferences 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.
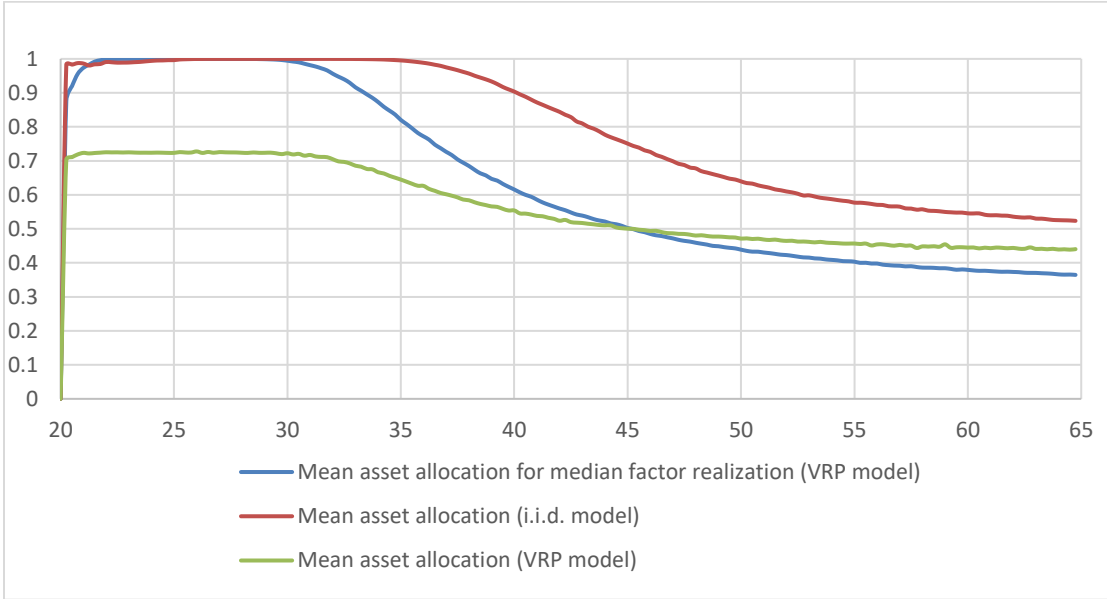
Figure 3 shows the expected portfolio return between the VRP model and i.i.d. model and their difference. The baseline preferences are Epstein-Zin with a risk aversion of 5 and an elasticity of substitution equal to 0.5 and a quarterly discount factor equal to 0.99. VRP is the variance risk premium model and the decision frequency is quarterly.
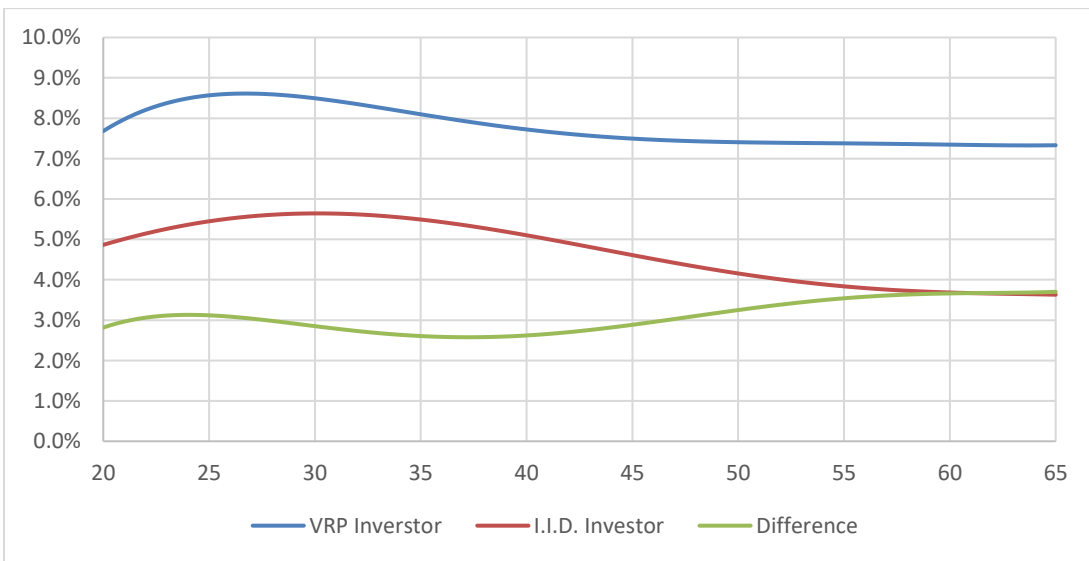
Figure 4 shows the mean share of wealth in stocks for the VRP and i.i.d. models and the target date funds (TDFs) that are constructed based on simulated shares of wealth in stocks and a multivariate regression on age and factor. In the i.i.d. model the factor state is irrelevant (as it should be). The 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.
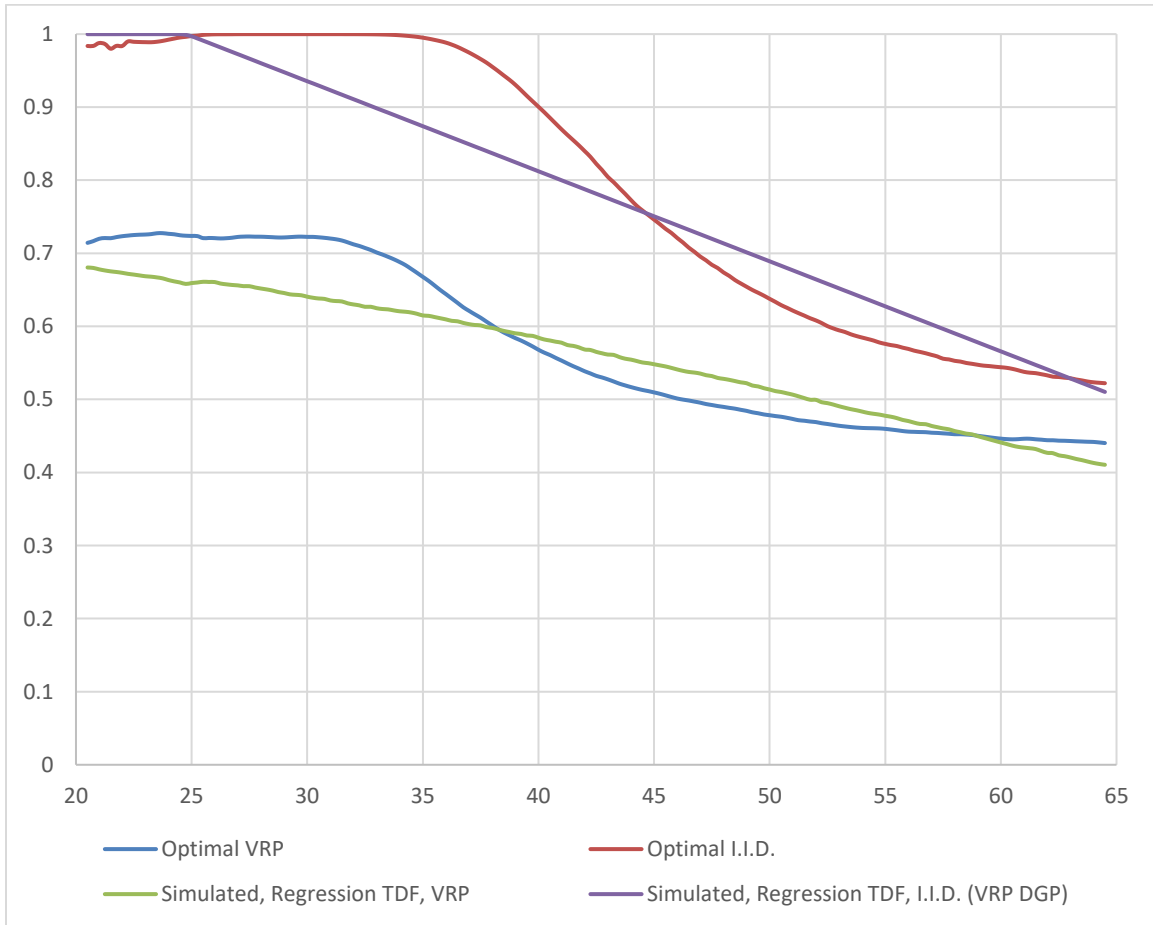
Figure 5 shows the share of wealth in stocks for the target date funds (TDFs) based on different factor realizations and the mean share of wealth in stocks for the VRP model. The data generating process (DGP) for stock returns in the simulation generating the simulated shares of wealth on which the TDF 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.