

LBS Research Online

[F Gomes](#), A Michaelides and Y Zhang
Tactical target date funds
Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/1585/>

[Gomes, F](#), Michaelides, A and Zhang, Y
(2022)

Tactical target date funds.

Management Science, 68 (4). pp. 3047-3070. ISSN 0025-1909

DOI: <https://doi.org/10.1287/mnsc.2021.3981>

INFORMS (Institute for Operations Research and Management Sciences)
<https://pubsonline.informs.org/doi/10.1287/mnsc.20...>

Users may download and/or print one copy of any article(s) in LBS Research Online for purposes of research and/or private study. Further distribution of the material, or use for any commercial gain, is not permitted.

Tactical Target Date Funds*

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

December 2020

Abstract

We propose target date funds modified to exploit stock return predictability driven by the variance risk premium. The portfolio rule of these tactical target date funds (TTDFs) is extremely simplified relative to the optimal one, making it easy to implement and to communicate to investors. We show that saving for retirement in TTDFs generates economically large welfare gains, even after we introduce turnover restrictions and transaction costs, and after taking into account parameter uncertainty. This predictability also appears to be uncorrelated with individual household risk, suggesting that households are in a prime position to exploit it.

JEL Classification: G11, G50

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

*We thank two anonymous referees, an associate editor and the editor (Tomasz Piskorski), as well as Jinhui Bai, Sylvain Catherine, Joao Cocco, Michael Haliassos, Anthony Lynch, Raimond Maurer, Philippe Mueller, Giovanna Nicodano, Lubos Pastor, Tarun Ramadorai, Paolo Sodini, Grigory Vilkov, Russ Wermers, Hao Zhou and seminar participants at the EFA and AFA 2020 Meetings, CEPR Household Finance Conference, Annual Conference on Macroeconomic Analysis and International Finance at the University of Crete, University of Cyprus, Cheung Kong Graduate School of Business, Durham University, FEP, Frankfurt School of Finance & Management, IAAE 2019, International 2019 Finance Conference in Moscow, PBC School of Finance (Tsinghua University), SFS Cavalcade, University of Lugano, University of Glasgow, and University of Luxembourg for comments and suggestions.

[†]Department of Finance, London Business School, London NW1 4SA, UK. E-mail: fgomes@london.edu.

[‡]Department of Finance, Imperial College Business School, 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.

1 Introduction

Conventional financial advice prompts households to invest a larger proportion of their financial wealth in the stock market when young and gradually reduce this exposure as they grow older. This advice is given by several financial planning consultants (for instance, Vanguard¹) who recommend target-date funds (TDFs) that reduce equity exposure as retirement approaches. The long term investment horizon of 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 context of life-cycle models with undiversifiable labor income risk (for example, Cocco, Gomes, and Maenhout (2005), Gomes and Michaelides (2005), Polkovnichenko (2007), and Dahlquist, Setty and Vestman (2018)).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 simple portfolio rules designed to capture time variation in expected returns can improve welfare for an investor saving for retirement.^{3,4} We explore three different popular return predictors: the dividend-price (DP) ratio, the CAY variable introduced by Lettau and Ludvigson (2001), and the variance risk premium (VRP) introduced by Bollerslev, Tauchen and Zhou (2009) and Bollerslev, Marrone, Xu, and Zhou (2014)). We first document that the VRP has the highest predictive power for future returns and, consistent with this, it generates economically large welfare gains in the context

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

²Benzoni, Collin-Duffresne, and Goldstein (2007), Lynch and Tan (2011) and Pastor and Stambaugh (2012) show that this conclusion can be reversed under certain conditions.

³In 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.

⁴The 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 (2017 and 2019) consider time variation in volatility, while Buraschi, Porchia and Trojani (2010) incorporate time-varying correlations.

of an optimal life-cycle portfolio choice model.⁵ Relative to an investor who assumes i.i.d. expected returns, an investor exploiting VRP predictability optimally earns a significantly higher expected return. This result holds even in the presence of 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 percentage points higher at each age (annually).

Crucially, we move beyond the optimal strategies implied by the model, and explore simpler strategies based on those optimal solutions that can be easily implemented by improved target date funds and can be easily explained 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 sceptical about complex financial products.⁶ It also follows the approach taken by the current TDFs, which do not use the exact policy functions of individual households, and instead offer an approximation that can be easily explained and implemented 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 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

⁵The welfare gains would potentially be even higher if we considered more recent predictors that have been shown to outperform the variance risk premium, such as the implied correlation or the correlation risk premium (see Buss, Schonleber and Vilkov (2018)).

⁶There 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.

⁷In a similar spirit, Dahlquist, Setty and Vestman (2018) study simple adjustments to the portfolio rules of TDFs to take this into account.

premium).⁸ We further truncate the fitted linear rule by imposing 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. Building on our initial discussion, we refer to those modified funds as Tactical Target Date Funds (hereafter TTDFs). Furthermore, we restrict the portfolio strategy of the TTDF by imposing a limit on its quarterly turnover.

When comparing the TTDF without turnover constraints with the TDF, we obtain a certainty equivalent gain of 5.72% for an investor with risk aversion of 5, and 10.1% for an investor with risk aversion of 10. As we impose turnover constraints, the certainty equivalent for the investor with risk aversion of 5 is still 2.37% and 1.35%, for a maximum rebalancing of 25 and 15 percentage points (pp) in the risky share, respectively. These values are economically large since they are comparable to the certainty equivalent gain from stock market participation, which is about 2.0%. In other words, even with a rebalancing limit of 25pp, the welfare gain obtained by switching from the TDF to the TTDF is comparable to the welfare gain obtained by becoming a stock market participant in the first place.

One concern with the previous calculations is that the welfare gains were computed in-sample. We address this concern in two ways. First, we estimate the predictive model in an initial sample (1990-1999) and use only that information to design the TTDF (more precisely, the TTDF with a tight turnover restriction). We then compare the real-time performance of this fund relative to the standard TDF over the subsequent period (2000-2016). This period was chosen even though the coefficients of the predictive regression are less stable exactly in the years immediately following our estimation window, before “recovering” in the final part of the sample. Nevertheless, we find that the TTDF would have outperformed the TDF, and that this improved performance is largely obtained by decreasing the magnitude of the losses in bad years (e.g. 2001, 2002 and 2008). These results highlight that the improved performance of the TTDF is not driven by excessive risk taking, on the contrary, it is often

⁸We also explore more sophisticated rules which naturally deliver higher wealth accumulation and utility gains but, for reasons just discussed relating to clarity and simplicity in communication, our baseline case remains the simple TTDF.

the result of a reduction in risk-taking in anticipation of lower expected returns.

Second, we explicitly model parameter uncertainty using a Bayesian approach. We find that the welfare gains from the new TTDF are almost unchanged relative to the original model. In the presence of parameter uncertainty the optimal portfolio rule is more conservative with the investor being more careful about exploiting potential predictability. Nevertheless, quantitatively, with quarterly rebalancing and short selling constraints, the optimal policy functions with parameter uncertainty are not substantially different from the behavior implied without parameter uncertainty. Moreover, the tight turnover restrictions further restrict portfolio changes. As a result, the impact of parameter uncertainty in designing the TTDF is small, and consequently our conclusions when comparing TTDF welfare gains relative to the TDF remain unchanged.

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. An additional extension considers an heteroskedastic model for returns where time-varying volatility is driven by the predictive factor, and the results are not significantly affected.

Naturally these tactical target date funds could be replicated by a combination of a standard target date fund and a predictability fund that uses the VRP strategy. But this would require an investor who not only has access to that second fund, but is also able to solve for the optimal weights across the two for each state-of-the world. In fact, the same argument can be made even more cleanly for the simple target date fund itself: it can be replicated by combining a pure index fund and cash. Moreover, in this simpler case the weights are only age dependent and therefore the strategy requires very limited financial knowledge to implement. To the extent that limited financial literacy and/or transaction costs (both financial and opportunity cost of time) have created such a large market for the simple TDF, the same forces should be even stronger for introducing the TTDF.

Our focus on the predictability driven by the VRP is also motivated by the high-frequency nature of this time variation in expected returns. The more traditional predictive variable, the dividend-yield, captures lower frequency movements (and is more persistent than the VRP) tending to be associated more closely with bad economic conditions and/or discount rate shocks, both of which affect households directly.⁹ We confirm this by documenting that excess return predictability from the dividend-yield is indeed associated with higher household consumption risk. This empirical result is also a direct implication of several equilibrium asset pricing models that endogenously generate these predictability patterns (e.g. Campbell and Cochrane (1999)). The existing portfolio choice literature is well aware of this and therefore carefully mentions that the results should not be interpreted as applying to a representative investor, but rather to an investor not exposed (or less exposed than the average) to such risks.¹⁰

In this paper we propose funds to be used by all investors so it is important that the average household is not exposed to these risks. Our underlying hypothesis is that VRP predictability is likely driven by constraints on banks, pension funds and mutual funds (e.g. capital constraints or tracking error constraints), unlikely to be significantly correlated with household-level risks. We provide supportive evidence for this argument using the Consumer Expenditure Survey (CEX). Specifically, we document that states of the world with high realizations of the VRP do not predict future decreases in household consumption growth, future increases in cross-sectional consumption risk, or decreases in household labor income growth. This evidence supports the hypothesis that households appear to be in a prime position to “take the other side” and exploit this premium. Furthermore, in general equilibrium, the fact that households own the financial intermediaries 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 keen to offset this by increasing the risk exposure in their individual portfolios.

The paper is organized as follows. Section II discusses the return predictors. In Sec-

⁹Bad economic conditions will tend to be associated with negative labor income shocks, while discount rate shocks might reflect increased household risk aversion.

¹⁰Michaelides and Zhang (2017) incorporate stock market predictability through the dividend-yield in the context of a life-cycle model of consumption and portfolio choice.

tion III we study the potential correlation between return predictability and household risk. Section IV outlines the life-cycle model and discusses the optimal policy functions and corresponding welfare gains. Section V discusses the design of the proposed TTDFs and their associated welfare gains, while Section VI incorporates parameter uncertainty and reports the corresponding out-of-sample performance. Section VII explores different extensions and robustness tests, and Section VIII provides concluding remarks.

2 Return Predictors and Stock Returns

2.1 VAR model for stock returns

Time variation in expected returns is captured by a predictive factor (f_t); following Campbell and Viceira (1999) and Pastor and Stambaugh (2012) we construct the VAR,

$$r_{t+1} - r_f = \alpha + \beta_f f_t + z_{t+1}, \quad (1)$$

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

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 bivariate normal variables with mean equal to zero and variances σ_z^2 and σ_ε^2 , respectively.

In our estimation we explore three different popular return predictors (f_t), the dividend-price (DP) ratio, the CAY variable introduced by Lettau and Ludvigson (2001), and the variance risk premium (VRP) introduced by Bollerslev, Tauchen and Zhou (2009) and Bollerslev, Marrone, Xu, and Zhou (2014)). The formulation allows for contemporaneous correlations between z_{t+1} and ε_{t+1} .¹¹

Our baseline comparison we will be a model with i.i.d. excess returns, in which case

$$r_{t+1} - r_f = \mu_r + z_{t+1}. \quad (3)$$

¹¹In the numerical solution of the model we approximate this VAR using Flodén (2008)'s variation of the Tauchen and Hussey (1991) procedure, designed to better handle the case of a very persistent AR(1) process. As discussed below, the CAY and the dividend-price ratio are very persistent variables.

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.

2.2 Return predictors

We use the CAY constructed by Lettau and Ludvigson (2001)¹² and we construct the DP from the value weighted returns of the S&P 500 index from the CRSP database. 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),

$$VRP_t \equiv IV_t - RV_t. \quad (4)$$

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 (2018).¹³ We convert the monthly realized variance to quarterly by simply 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), replicating and extending essentially the original Bollerslev, Tauchen and Zhou (2009) measure.

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 7.8%. Relative to the other two predictors the variance risk premium has a very high kurtosis and negative skewness. In terms of volatility, the CAY variable has the highest standard deviation. Of course, these results do not necessarily translate into the implied forecasts of expected returns, since those will also depend on the coefficients of the corresponding VAR. Finally, while the CAY and DP are highly persistent, with a first-order autocorrelation of 0.93 and 0.82 respectively, the VRP is even slightly negatively correlated (-0.18 first-order autocorrelation).

¹²This can be found at Martin Lettau's website: <https://sites.google.com/view/martinlettau/data>.

¹³Available at <https://sites.google.com/site/haozhouspersonalhomepage/>.

2.3 VAR Estimation

Table 2 reports the estimation results for the VAR model (equations (1) and (2)) when using the three different predictive factors. To facilitate the discussion we use the terminology VRP model, CAY model and DP model, when referring to the VAR models that use VRP, CAY, and DP, respectively, as the predictive factor.

The CAY and DP models have an R^2 of 5.4% and 6.2%, respectively, compared to 15% for the VRP model. In addition, the statistical significance of β_f is also higher in the VRP model, with a t-statistic of 4.48, versus 1.40 and 2.83 for CAY and DP, respectively. These results suggest that the variance risk premium will be the most effective predictor to consider in the portfolio choice model. Our quantitative estimates for the VRP model are also 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

$$\sigma_z^2 = \text{Var}(r_t) - \beta_f^2 \sigma_f^2. \quad (5)$$

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) this is a large negative number (see, for example, Campbell and Viceira (1999) and Pastor and Stambaugh (2012)). Here, this correlation is estimated as slightly negative over the whole sample for the VRP. Nevertheless, in Table 8 we show that the correlation is strongly negative for the subsamples that exclude the financial crisis.¹⁴

Following the life-cycle portfolio choice literature, we adjust our return process in all specifications to deliver an unconditional equity premium below the historical average, namely 4% at an annual frequency. The net constant real interest rate, r_f , is set at 0.37% corresponding to 1.5% at an annual frequency.

¹⁴As discussed in Section 6.3, the performance of the TTDF remains high even during these periods.

2.4 Return predictors and return volatility

The model presented in section 2.1 assumes homoskedastic returns. This is a common assumption in the life-cycle portfolio choice literature to avoid having to introduce lagged return volatility as an additional state variable in the model.¹⁵ Since return heteroskedasticity would affect both the model without first-moment predictability and the models with the predictive factors, we follow the previous literature and consider a model with constant variance. In our context this assumption could be problematic if time-variation in return volatility is associated with a specific predictive factor. More precisely, if $t + 1$ -volatility is expected to be high when a given factor predicts high expected returns at $t + 1$, then we will be over-stating the welfare gains by ignoring this link.¹⁶ We address this concern by estimating the following equation

$$Var_{t+1} = a + bf_t + v_{t+1} \quad (6)$$

where f_t is any of the three predictive factors. The results are shown in Table 2, Panel B.

We find that the variance risk premium is not a statistically significant predictor of future volatility at a quarterly frequency. This result validates our assumption that time-variation in expected returns, as predicted by the VRP, is not associated with changes in the second moment of returns. It is interesting to note, however, that this result arises because we are using a quarterly frequency. As also shown in Table 2, Panel B, when we use monthly data in our estimation we find a statistically significant correlation between current VRP and future stock return volatility. Despite having estimated a non-significant coefficient, as a robustness exercise we later consider a version of the model with equation (6), and we use our estimate of b as the calibrated coefficient.

Considering the other two predictors, CAY is also statistically uncorrelated with the future volatility of stock returns, while for the dividend-price ratio there is a statistically significant negative coefficient.

¹⁵In addition these models are usually solved at an annual frequency, and heteroskedasticity in returns is much weaker in that frequency.

¹⁶By the same logic, if $t+1$ -volatility is actually expected to be lower, then we will actually be under-stating the welfare gains.

3 Return Predictors and Household Risk

3.1 Discussion

The optimality of increasing the allocation to stocks when expected returns are high will be over-stated if this is accompanied by an increase in risk for households. We have already documented a high value of the VRP does not predict a higher variance of stock returns, but it could be associated with other economic risks that impact households directly. Therefore, it becomes important for our analysis that this is not the case, and in this section we provide supporting evidence for this argument.

It is important to clarify that we are not arguing that the changes in expected returns, as forecasted by the VRP, do not reflect risk. 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.¹⁷ 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.¹⁸ 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 Data and Variable Construction

We use non-durable consumption and services from the Consumer Expenditure Survey (CEX).¹⁹ We exclude durables, implicitly assuming that utility is separable between durables

¹⁷For example, tracking error constraints for mutual funds or VAR constraints for banks.

¹⁸Naturally, 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.

¹⁹Our internet appendix provides further details on data construction.

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 data construction follows the one in Malloy et. al. (2009). We first drop non-urban households, households residing in student housing, and households with incomplete income responses. We also exclude household-quarters in which a household reports nonzero consumption for more than 3 or less than 3 months in any one interview, or where consumption is negative. Likewise, we remove observations with extreme values, namely those with consumption growth above percentile 97.5, and below percentile 2.5. 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”. In our data the stock market participation rate is around 19%, which is similar to the rate reported by Malloy et. al. (2009) for the earlier version of this sample.

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 consumption growth measure.

Our model applies primarily to stockholders; since stockholders face different risks from non-stockholders, we estimate separate regressions for the two groups. We compute the average consumption growth rate for a particular group (for instance, stockholders) for different horizons $s=1, 2, 4,$ and 8 by averaging the log consumption growth rates as

$$\frac{1}{N} \sum_{i=1}^N [c_{i,t+s} - c_{i,t}], \quad (7)$$

where $c_{i,t}$ is the quarterly log consumption of household i at time t .

As discussed below, we also investigate whether our measure of expected returns is associated with an increase in cross-sectional consumption risk. For this purpose we compute the higher-order moments of the cross-sectional distribution of consumption, namely the standard deviation, skewness and kurtosis. Since higher order moments are not additive, unlike the mean consumption growth rate, we can only compute them for $s = 1$ and $s = 2$, since they can be constructed directly for the same group of households.

3.3 Consumption risk

We start by considering whether higher expected returns are associated with lower expected future consumption growth, by estimating the following regressions:

$$\frac{1}{N} \sum_{i=1}^N [c_{i,t+s} - c_{i,t}] = \alpha_c^s + \beta_c^s \cdot f_t^j + \varepsilon_t^c, \quad s = 1, 2, 4, 8, \quad j = VRP, CAY, DP, \quad (8)$$

where f_t^j denotes the realization of the predictive factor j at time t .

The estimates of β_c^s are shown in Panel A of Table 3 for the sub-sample of stockholders.²⁰ The standard errors are computed using a Newey-West estimator that allows for autocorrelation of up to $s - 1$ lags when $s > 1$. The third and fourth rows in Table 3, Panel A, show that β_c^1 is non-significant for all values of s , when the predictor variable is the VRP. In fact, most point estimates are even positive. This confirms that a high expected return, as predicted by the variance risk premium, is not associated with lower expected future consumption growth rate, either in the next quarter ($s = 1$), or even as far as 2 years ahead ($s = 8$). It is still possible, however, that the higher expected returns predicted by the variance risk premium are associated with an increased dispersion of future consumption growth. In Panel B we further exploit the cross-sectional dimension of the CEX to repeat the previous regressions for higher-order moments of the cross sectional distribution of consumption growth, namely the standard deviation, skewness and kurtosis.²¹ The results in the third and fourth rows of

²⁰The results for non-stockholders are presented in the internet appendix. They deliver the same qualitative conclusions, although we find less statistically significant coefficients.

²¹As previously discussed these moments can only be constructed for $s = 1$ and $s = 2$, so we can only

Panel B show that the VRP is not associated with an increase in any of these three measures of cross-sectional consumption risk, both for $s = 1$ and for $s = 2$.

When the predictor variable is the dividend-yield, the point estimates in Table 3, Panel A, are negative and statistically significant for all horizons greater than $s = 1$. Therefore, unlike the VRP, high values of the dividend-yield also predict lower future expected consumption growth. In other words, if an investor increases her risky share when the dividend-yield is high, she will be exposed to more return risk exactly when her consumption is expected to fall. This result is in fact implied by a large asset pricing literature that rationalizes the dividend-yield predictability in equilibrium (e.g. Campbell and Cochrane (1999)).

Interestingly, for the CAY we find the same results as for the VRP. None of the coefficients in Panel A of Table 3 is statistically different from zero, so there is no statistically significant correlation between higher expected returns, as predicted by the CAY, and future expected consumption growth. Turning to Panel B of Table 3 we find that high values of CAY are however associated with an increased volatility of future consumption growth. Comparing with the VRP, the overall results are not substantially different though, suggesting that the CAY could be another interesting predictor to consider in designing TTDFs.

3.4 Labor income risk

Our previous results imply that, even if the return predictability from the VRP is associated with increased economic risk, households are either not directly affected by those risks, or able to smooth them, and therefore not reflected in their consumption. However, given the important role of labor income in our portfolio choice model we also explore whether return predictability is associated with higher labor income risk.

More precisely, we follow the same methodology used for estimating the impact on consumption risk, by estimating the following series of regressions²²

$$\frac{1}{N} \sum_{i=1}^N [y_{i,t+s} - y_{i,t}] = \alpha_y^s + \beta_y^s \cdot f_t^j + \varepsilon_t^y, \quad s = 1, 2, 4, 8, j = VRP, CAY, DP, \quad (9)$$

report regression results for these two horizons.

²²We again use CEX data because it has a quarterly frequency, which is not the case with PSID data, for example.

where f_t^j again denotes the realization of the predictive factor j at time t .

The results are shown in Table 3, Panel C, and we again find no statistically significant coefficients for the VRP or CAY. Most of the coefficients for the VRP are even positive, although not statistically significant. For the dividend-price ratio there is again evidence of increased future risk, with a negative statistically significant coefficient at eight quarters.

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. We 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),

$$V_a = \max \left\{ (1 - \beta)C_a^{1-1/\psi} + \beta (p_a E_a(V_{a+1}^{1-\gamma}))^{\frac{1-1/\psi}{1-\gamma}} \right\}^{\frac{1}{1-1/\psi}}, \quad (10)$$

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

$$W_{a+1} = [\alpha_a R_{t+1} + (1 - \alpha_a) R_f](W_a - C_a) + (1 - h_{a+1}) Y_{a+1} \quad (11)$$

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

$$\alpha_a \in [0, 1]. \quad (12)$$

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 by²³

$$Y_a = \exp(g(a)) Y_a^p U_a, \quad (13)$$

$$Y_a^p = Y_{a-1}^p N_a \quad (14)$$

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

²³We 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.

to the expected and unexpected stock returns.

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

$$w_{a+1} = \frac{1}{N_{a+1}} [R_{t+1}\alpha_a + R_f(1 - \alpha_a)](w_a - c_a) + (1 - h_{a+1}) \exp(g(a + 1))U_{a+1}. \quad (15)$$

As common in the literature the retirement date ($K = 65$) is exogenous, and retirement income is a deterministic function of working-life permanent income

$$Y_a = \lambda Y_K^p \text{ for } a > K, \quad (16)$$

where λ is the replacement ratio.

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 consider different values for the preference parameters, to explore how the welfare gains might change across investors. More precisely we consider risk aversion coefficients of 2, 5 and 10, discount factors of 0.9875 and 0.995 (annual equivalents of around five percent and two percent, respectively), and elasticities of intertemporal substitution of 0.5 and 1.5.

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. In Figure 2.1, 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 computed using the corresponding policy functions.

The portfolio share from the i.i.d. model follows the classical hump-shape pattern (e.g. Cocco, Gomes and Maenhout (2005)).²⁴ 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

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

²⁴The increasing pattern early in life is barely noticeable because under our calibration the average optimal share at young ages is (already) close to one.

Even though under the two scenarios the expected return on stocks is the same, Figure 2.1 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.²⁵

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 2.1 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

$$E[\alpha_a(f)] < \alpha_a[E(f)] \text{ for } a < 45 \quad (18)$$

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 $\alpha_a[E(f)]$. In some cases, depending on the volatility of the predictive factor, the expected next period stock return becomes negative, and the optimal share of wealth in stocks jumps to zero. As a result, the optimal average allocation of the VRP investor is sometimes far below $\alpha_a[E(f)]$.

Building on the previous intuition, it is not surprising to find that the inequality sign 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:

$$E[\alpha_a(f)] > \alpha_a[E(f)] \text{ for } a > 45 \quad (19)$$

This comparison suggests that the welfare gains from the VRP model are likely to be much higher if we relax the short-selling constraints, which motivates our discussion of this par-

²⁵The 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.

ticular extension in Section 7.

Combining inequalities (17) and (18) it is easy to see that, until age 45, we have:

$$E[\alpha_a(f)] < E[\alpha_a^{iid}], \quad (20)$$

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 (18) is replaced by (19), although the difference between the optimal allocation of the VRP and i.i.d. investors decreases, equation (20) still holds: the i.i.d. investor has a higher average allocation than the VRP investor.

In Figure 2.2 we plot three different moments of the risky share under the VRP model: the mean (already in Figure 2.1), and plus one and minus one standard deviation based on simulated portfolios. Figure 2.2 illustrates that the optimal portfolio share from the VRP model implies a very high level of turnover. For this reason, when we later design the portfolio rules for the tactical target date funds, we will explicitly impose constraints on maximum portfolio rebalancing.

4.5 Portfolio returns

We now study the differences in expected returns between the VRP and i.i.d. investors, assuming they start with zero initial financial wealth and they face the same labor income realizations. In Figure 3.1 we plot the (annualized) average expected portfolio returns

$$E(R_{t+1}^P) = \alpha_a E_t[R_{t+1}] + (1 - \alpha_a)R_f, \quad a = 1, \dots, T, \quad (21)$$

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 ($\bar{\alpha}$), this would be a flat line. For example, if $\bar{\alpha} = 1$, 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 2.1 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 percentage points. This difference is therefore at its maximum exactly when these investors have the highest wealth accumulation.

Figure 3.2 shows the range of possible outcomes we might expect by plotting the expected return for the VRP and i.i.d. investor over the life-cycle, but also with a plus one, and minus one, standard deviation band. At plus one standard deviation, the two models deliver similar returns, particularly early in life. This is expected since, at these ages, the optimal allocation to stocks at the mean expected return is 100%, or very close to it. As a result, if the VRP model predicts higher expected returns, the investor cannot increase her allocation further and is therefore unable to exploit this predictability. However, this constraint does not prevent the investor from decreasing her allocation when expected returns are below average. For that reason the VRP model significantly outperforms the i.i.d. model when in bad states of the world, as shown by the minus one standard deviation lines. Overall, the results in Figure 3.2 show that the improved performance of the VRP is primarily obtained by minimizing the impact of low returns.

4.6 Welfare metrics and welfare gains

4.6.1 Welfare metrics

Our welfare metrics are consumption certainty equivalent gains (CEG), defined as the percentage increase in life-time certainty equivalent consumption

$$CEG_{20}^i = \frac{CEQ_{20}^i}{CEQ_{20}^{iid}} - 1, \quad i = VRP, CAY, DP, \quad (22)$$

where CEQ_{20}^i is the life-time (from age 20) certainty equivalent consumption obtained from the value function with the optimal policy rules implied by model i , and CEQ_{20}^{iid} is the lifetime certainty equivalent consumption obtained with the policy rules implied by the i.i.d. model. By construction, the consumption certainty equivalent gains take into account both differences in expected returns/wealth accumulation and differences in risk.²⁶

Consistent with the focus of our paper to design improved target date funds for retirement savings, we also report welfare gains at age 65, i.e. consumption certainty equivalent gains measured using the age-65 value functions,

$$CEG_{65}^i(W_{65}^i) = \frac{CEQ_{65}^i(W_{65}^i)}{CEQ_{65}^{iid}(W_{65}^{iid})} - 1, \quad i = VRP, CAY, DP, \quad (23)$$

where CEQ_{65}^i and CEQ_{65}^{iid} are the consumption certainty equivalents when considering the retirement period only (from age 65), and evaluated at the average age-65 wealth level implied by the model.²⁷ To facilitate comparisons across models, we also compute this certainty equivalent using the same wealth level for all, namely the average age-65 wealth accumulation of the VRP model:

$$CEG_{65}^i(W_{65}^{VRP}) = \frac{CEQ_{65}^i(W_{65}^{VRP})}{CEQ_{65}^{iid}(W_{65}^{iid})} - 1, \quad i = VRP, CAY, DP. \quad (24)$$

Finally, we also report age-65 certainty equivalent gains obtained by imposing an identical

²⁶Furthermore, risk is evaluated beyond just considering the second moment of returns. Since we use an Epstein-Zin utility function, the investor cares not only about first and second moments, but about the full distribution of returns.

²⁷For simplicity we omitted wealth levels in equation (22) since we assume that all agents start with zero wealth at age-20.

pre-retirement consumption behavior,²⁸

$$CEG_{65}^i(W_{65}^i(C_{20-65}^{iid})) = \frac{CEQ_{65}^i(W_{65}^i(C_{20-65}^{iid}))}{CEQ_{65}^{iid}(W_{65}^{iid})} - 1, \quad i = VRP, CAY, DP \quad (25)$$

where $W_{65}^i(C_{20-65}^{iid})$ is the age-65 wealth obtained when using, from age 20 to age 65, the portfolio rules from model i and the consumption rules from the i.i.d model.²⁹ This calculation equalizes pre-retirement consumption behavior, and therefore takes the view of two otherwise identically-behaving investors who simply invest in different portfolios, and therefore any differences in consumption are exclusively determined by the differences in the realized returns on those portfolios. By comparing (23) and (25) we can isolate the role of the differences in pre-retirement consumption behavior for determining the age-65 welfare gain.

4.6.2 Results

The results for the 3 different models are shown in Table 4. Panel A reports results for different values of the coefficient of relative risk aversion, 2, 5 and 10, for fixed values of the EIS (0.5) and the discount factor (0.9875 quarterly). In Panel B we report results for alternative values of the EIS (1.5) and the discount factor (0.995 quarterly), while keeping the coefficient of relative risk aversion constant at 5.

In all specifications the welfare gains increase with risk aversion. Wealthier individuals have more to gain by increasing the expected return on their wealth, and in life-cycle models with undiversifiable income risk a higher risk aversion typically leads to higher wealth accumulation. For the same reason, the welfare gains are higher when measured at age-65, since agents are wealthier around retirement age.

Crucially, for all three values of risk aversion, and regardless of the welfare metric, there is a clear ranking of the models: the VRP model yields the highest welfare gain, followed by the CAY model, and then the DP model. The ranking remains unchanged even when we compare the models at the same level of wealth at age-65 ($CEG_{65}^i(W_{65}^{VRP})$) or at age-20

²⁸More specifically, we assume that before retirement the agent uses the consumption functions of the i.i.d. model in all cases.

²⁹By construction, $W_{65}^{iid} = W_{65}^{iid}(C_{20-65}^{iid})$.

(CEG_{20}^i) .³⁰ Based on these results, we take the VRP model as the basis for the tactical target-date funds that we will study in the remainder of the paper.

The age-20 welfare gains are smaller because all future gains are discounted, but they are economically very large, especially for the VRP model. For comparison, the corresponding certainty equivalent gains from stock market participation obtained in life-cycle models are in the order of 1% to 3% depending on the calibration (e.g. Cocco, Gomes and Maenhout (2005)).³¹ In other words, the benefit from moving from the i.i.d. model to VRP model is larger than the benefit obtained from becoming a stockholder in the first place.

In Panel B, we report results for alternative values of the EIS (1.5) and the discount factor (0.995 quarterly), while keeping the coefficient of relative risk aversion constant at 5. As we increase the EIS, or the discount factor, the welfare gains are even larger than under our baseline calibration and the intuition is the same as when comparing risk aversion in Panel A. These two cases lead to higher wealth accumulation, and wealthier investors benefit more from improving the return on their wealth. As in Panel A, we again conclude that the VRP model delivers the highest welfare gains for all combinations of parameter values.

5 Tactical Target-Date Funds

Having concluded that the VRP is the best expected return predictor, we now proceed to incorporate this factor in an improved target-date fund. 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 implemented 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,

³⁰Since agents start at age-20 with zero wealth, the age-20 certainty equivalent gain is by definition evaluated at the same value of cash-on-hand in all cases.

³¹For the baseline calibration with a risk aversion coefficient of 5, the certainty equivalent gain from participating in the stock market (under the i.i.d. model) is 1.98%.

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 with possibly limited financial literacy, and have to make the final decision on where to allocate their retirement savings.

In the same spirit, in our baseline specification we derive a straightforward portfolio rule that can be implemented by a tactical target date fund (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 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 mis-specification. Finally, in this section, we impose short-selling constraints on both the TDF and the TTDF. Later on we discuss the results obtained when we relax these constraints.

5.1 Designing Tactical Target-Date Funds

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

$$\alpha_{iat} = \theta_0 + \theta_1 \cdot a + \theta_2 \cdot f_t + \varepsilon_{iat}. \quad (26)$$

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 (that affects portfolios linearly through (θ_2)). However, as previously argued, this is simple to implement and easier to explain to investors. Since these linear rules do not satisfy the original short-selling

constraints, we impose those again ex-post, so

$$\alpha_{iat}^{TTDF} = \text{Max}(\text{Min}(\theta_0 + \theta_1 \cdot a + \theta_2 \cdot f_t, 1.0), 0.0) \quad (27)$$

Table 5 reports the regression results from these rules for our baseline cases of relative risk aversion equal to 5 and 10, EIS equal to 0.5 and discount factor equal to 0.9875.³² For comparison, we also report the results for the i.i.d. model. As we increase risk aversion the average equity exposure decreases, reflected in a lower value of θ_0 . Interestingly, the more risk averse investor also exploits time-variation in expected returns less (lower θ_2).

These results are also shown graphically in Figure 4 for the average share of wealth in stocks over the life cycle for the baseline cases. As previously explained, in the VRP world the investor moves more aggressively from positive to sometimes zero investment positions and this explains the lower average share of wealth in stocks relative to the i.i.d. model. This behavior is reflected in the design of the mutual fund associated with each model. The TTDF (TDF) associated with the VRP (i.i.d.) model is drawn based on a linear regression of all simulated portfolios on the factor and age. For simplicity, and for comparison purposes, we show the linear rule by averaging over all factors for the TTDF: this predicted share of wealth in stocks is a straight line across the average share of wealth in stocks generated by the VRP model. On the other hand, the effect of the factor is irrelevant for the TDF because the TDF is based on the i.i.d. model. Figure 4 therefore shows in a parsimonious way the average differences between the TTDF and TDF design.

5.2 Turnover Restrictions

One potential concern with the TTDFs is that their implementation might imply a very high portfolio turnover. Indeed, in our simulations the average (annualized) portfolio turnover implied by the TTDF is 209% indicating that tactical asset allocation implies a very volatile asset allocation behavior over the life cycle. By comparison, the average turnover of the

³²As shown in the previous section, the welfare gains would be larger (smaller) if we considered a higher (lower) risk aversion, EIS or discount factor.

typical mutual fund is 78% (see Sialms, Starks and Zhang (2013)).³³ Therefore, we further restrict the portfolio strategy of the TTDF by imposing an explicit turnover restriction. 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

$$\alpha_a = \begin{cases} \alpha_{a-1} + k & \text{if } \alpha_a^* > \alpha_{a-1} + k \\ \alpha_a^* & \text{if } |\alpha_a^* - \alpha_{a-1}| < k \\ \alpha_{a-1} - k & \text{if } \alpha_a^* < \alpha_{a-1} - k \end{cases} \quad (28)$$

where α_a^* is the optimal allocation in the absence of the constraint. In our analysis, we consider two thresholds ($k = 0.25$ and 0.15), in addition to the unconstrained case ($k = 1$).

We impose equation (28) in two steps. First, we impose it in an extended version of the dynamic programming problem where we add the lagged portfolio choice (α_{a-1}) as an additional state variable. More precisely, we re-formulate the optimization problem with the added constraint (28) and the additional state variable. This guarantees that the constraint holds under the optimal policies. However, since the TTDF rules are derived from linear regressions (26), they are an approximation to the optimal rules and therefore might not satisfy the original constraint, in the same way that they might not satisfy the short-selling constraint and we have to impose that ex-post as well.³⁴ Therefore, in the second step, we impose constraint (28) directly on the estimated TTDF rules.³⁵

Figure 5 illustrates the impact of these turnover restrictions. It shows the life-cycle portfolio allocation of both the unconstrained TTDF and the TTDF with the $k = 0.15$ turnover restriction. The figure plots the allocation for different realizations of the predictive factor (f), namely its mean (0.494%) and 1.14 standard deviations above and below the mean, respectively.³⁶ As we can see, in the absence of any restrictions the TTDF allocation changes

³³For the standard TDF (i.e. the one that replicates the optimal allocation of the i.i.d. investor) the average turnover is 23%.

³⁴Naturally, we estimate the regressions again as we change k , but still there is no guarantee that the fitted rules will satisfy the original constraint.

³⁵From an implementation perspective this is still a transparent rule that is 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 intuitive constraint.

³⁶Those apparently arbitrary values correspond to actual points on the grid for the state space, while plus and minus 1 standard deviations do not.

by almost 40% for an approximate one standard deviation movement in the predictive factor. Even in the presence of short-selling constraints this creates the very large turnover numbers that we have reported. In contrast, when we impose the $k = 0.15$ constraint the share invested in equities is much less volatile. For a 1 standard deviation movement in the VRP, the average change in the risky share is now about 10 pp.

5.3 Utility gains from Tactical Target Date Funds

5.3.1 Welfare Metrics

Having identified a feasible portfolio rule for the TTDF we now proceed to compute the corresponding certainty-equivalent utility gains. We report age-20 certainty equivalent gains (CEG_{20}^i , equation (22)), and the age-65 certainty equivalent gains computed for an identical pre-retirement consumption rule ($CEG_{65}^i(W_{65}^i(C_{20-65}^{iid}))$, equation (25)).³⁷ In comparing different rules we assume the same asset allocation rules *after retirement*, that is, 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 either used the optimal consumption rules, or allowed the investor to exploit time-variation in the risk premium during retirement as well; 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.

5.3.2 Results

Table 6 reports the welfare gains for the TTDF rule (equations (26) and (28)) for different values of risk aversion, 5 and 10, and different turnover limits, 15%, 25% and 100%). With a maximum rebalancing limit of 0.25, the average annual turnover of the fund falls from 209% and 181% to around 103% and 97.4%, respectively for risk aversion of 5 and 10. When the limit is even stricter (0.15), the mean turnover for the two funds is now only 66.2% and 62.8%, which is even below that of the typical mutual fund (78% as mentioned above). In

³⁷Results for the other 3 welfare metrics yield similar conclusions and are available upon request.

addition we also consider different values of the *additional* fund fee (Extra Fee) charged by the TTDF relative to the fee of a standard TDF: 0, 40 and 100 basis points, in annualized values. This extra fee might capture higher expenses as a result of implementing the market strategy, or simply higher profits for the mutual fund provider.

When the TTDF charges the same fee as the TDF (Extra Fee= 0), the expected increase in age-65 wealth accumulation for an investor with risk aversion of 5 is 183%, 40% and 8.99%, for $k = 1$, $k = 0.25$ and $k = 0.15$, respectively. For the unconstrained fund ($k = 1$) the cross-sectional standard deviation of (age-65) wealth is also much higher (247%), but nevertheless the investor is substantially better off, even in risk-adjusted terms, with an age-20 certainty equivalent gain of 5.72%. For comparison, the welfare gains from following the optimal portfolio rule without any turnover restrictions was 8.18% (Table 4). The TTDF's ability to capture almost 70% of the optimal gain is impressive since it is based on a simple linear approximation (equation (27)) and, furthermore, it only applies to investments before retirement age, so that the investor does not exploit predictability after age 65, unlike in the previous results (Table 4).³⁸

As we consider the cases with the tight turnover constraint the increase in wealth accumulation is less striking (40% and 8.99%), but there is also much smaller increase in the volatility of wealth.³⁹ As a result, the certainty equivalent gains are still economically large: 2.37% and 1.35%, respectively for $k = 0.25$ and $k = 0.15$. These values increase to 4.84% and 2.84% respectively, for the investor with risk aversion equal to 10. Even if we assume a higher fee for the TTDF, the age-20 certainty equivalent gains remain economically significant in almost all cases. For example, for an investor with risk aversion of 10, the unconstrained fund still yields a certainty equivalent gain of 7.95% when charging an extra fee of 100 basis points per year. With the strict turnover limits the corresponding age-20 welfare gain are still 2.85% and 0.92%, which are of comparable magnitude to the value of investing in the stock market under the i.i.d. scenario (2%).

Overall, the results in Table 6 confirm that it is possible to design a relatively simple

³⁸Below we consider an extension where the TTDF also includes the retirement period and welfare gains are naturally even larger.

³⁹For the case with $k = 15\%$ there is even a smaller reduction in the volatility of age-65 wealth.

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.

Finally, we also compute the maximum annual fee that the TTDF could charge to its investors, i.e. the fee that would set their age-20 certainty equivalent to zero. For an investor with risk aversion of 5, that fee is 470, 213 and 141 basis points, respectively for the three turnover limit cases (100%, 25% and 15%). These numbers increase to 522, 226, and 163 basis points when considering the fund for the investor with risk aversion of 10.

6 Parameter Uncertainty

One concern with the previous calculations might arise from the welfare gains being computed ignoring parameter uncertainty. In this section we address this concern in two ways. First, we incorporate parameter uncertainty in a Bayesian framework (e.g. Barberis (2000)). Second, we estimate the predictive model in an initial sample (1990-1999) and evaluate the performance of the TTDF over a subsequent period (2000-2016).

6.1 Bayesian approach

In this section we take into account parameter uncertainty using a standard Bayesian approach. For computational reasons we only consider parameter uncertainty over the two more important parameters: β_f , the predictive coefficient in the expected return equation, and ϕ , the persistence of the factor. We assume that the posteriors over the two parameters are independent and, under the assumption of diffuse priors, are given by

$$\beta_f \sim N(\widehat{\beta}_f, \widehat{\sigma}_{\beta_f}) \text{ and } \phi \sim N(\widehat{\phi}, \widehat{\sigma}_{\phi}),$$

where $\widehat{\beta}_f$ and $\widehat{\phi}$ are the corresponding point estimates (3.50 and -0.18 , respectively), and $\widehat{\sigma}_{\beta_f}$ and $\widehat{\sigma}_{\phi}$ are the standard errors from the estimation (0.78 and 0.09, respectively). We approximate both posterior distributions using standard Gaussian quadrature methods, just as for the other random variables in the model. Using the optimal solution from the model

with parameter uncertainty we repeat the previous process. We first fit a new TTDF rule and then evaluate the corresponding welfare gains relative to the TDF rule.

The results are reported in Table 7 for the baseline risk aversion coefficient of 5. The numbers are very similar to their counterparts in Table 6, when we did not consider parameter uncertainty. Our conclusions are therefore robust to parameter uncertainty concerns.

One explanation for this result is the inclusion of constraints in our baseline TTDF rule, namely the short-selling constraints and the turnover restrictions. In the absence of any constraints, parameter uncertainty will make the investor follow a more conservative portfolio rule, i.e. a portfolio rule closer to the one implied by the i.i.d. model. However, the presence of constraints, and in particular the tight turnover restrictions, have exactly the same effect. Therefore, if the investor is already more constrained by the turnover limit than what parameter uncertainty would imply, then adding parameter uncertainty will not meaningfully change the results, and this is indeed what we find.⁴⁰

Even if the policy functions are largely unaffected, parameter uncertainty introduces a second potential source of welfare loss by increasing the ex-ante standard deviation of expected future outcomes. When the coefficient on the predictive regression is not very precisely estimated, this effect can be significant (Pastor and Stambaugh (2012)). However, as shown in Table 2, the predictive coefficient on the VRP is very precisely estimated (with a t-statistic of 4.48), therefore this effect is small in our case.

6.2 Predictive model in different sub-samples

In Table 8 we report the VAR estimates for three different sub-samples: 1990-1999, 2000-2009 and 2010-2016. We see that the 3 different periods yield estimates that are broadly similar but also with some non-trivial differences. For example, even though the coefficient on the predictive regression is always positive, it falls to 2.01 in the middle period, compared with 5.85 and 5.23 in the other two periods. In the full sample estimation (Table 2), the correlation between realized returns and expected returns was only marginally negative. Here, we can see

⁴⁰Consider a hypothetical simplified example where, given the current realization of the factor, the unconstrained investor would like to increase her allocation to 80% in the absence of parameter uncertainty and to 75% in the presence of parameter uncertainty. If the maximum turnover constraint already implies that she cannot increase her allocation above 70%, then both portfolio rules would give the same allocation.

that this is the result of combining a positive correlation in the middle period with negative correlations in the other two. The results in Table 8 suggest that the TTDF's performance is particularly at risk right after our estimation window, and this further motivates our choice of this window for the out-of-sample exercise in the next sub-section.

6.3 Out-of-sample returns and wealth accumulation

We now repeat the design of the TTDF but based on the VAR estimated for the period 1990 to 1999. More precisely, we solve and simulate the model again using the data generating process from this VAR, and then use the simulated policy functions to estimate equation (26) again. For computational reasons we keep the TTDF rule constant throughout the exercise and do not update the model every quarter. This lowers the out-of-sample performance of the TTDF. Moreover, we restrict the turnover of the fund with the tightest value of this constraint, i.e. $k = 0.15$.

Figures 6.1 and 6.2 show the results for the 20-year-old investor. We report annual returns to facilitate the exposition but, as before, they are based on an underlying quarterly model and quarterly simulation. Figure 6.1 reports the cumulative returns over the period 2000-2016 for both the TTDF and the TDF. We can see that the TTDF out-performs from early on, and the gap between the two funds increases over time. By the end of the period the investor choosing the TTDF has accumulated 24% more wealth than the other investor. Figure 6.2 shows the (annualized) period-by-period returns and provides a better understanding of this superior performance. In good periods, the TTDF actually tends to under-perform (e.g. 2005 and 2012-2014). It is in bad times that the TTDF does consistently better, namely in the years 2001, 2002 and 2008. In these years the market timing investor is able to mitigate her losses relative to the TDF investor. Returning to figure 6.1 we can indeed confirm that the performance of two funds starts to diverge around 2001 and this difference increases again around 2008.

We can understand the superior performance of the TTDF in bad times from Figure 4, where we see that the average allocations of these funds are very high for young investors. As a result, in the presence of short-selling constraints, these investors can benefit more

from decreasing their equity exposure in bad times than from increasing it in good times. It is easier for them to hit the 100% constraint when trying to exploit high expected returns than to hit the 0% constraint when facing low expected returns. To further illustrate this intuition, Figures 7.1 and 7.2 show the results for a 50-year-old investor. As we saw in Figure 4, while the average risky share of the 20-year old investor is close to 70%, for the 50-year old investor this number is almost exactly 50%. As a result, the benefits from the TTDF are now more evenly distributed across booms and busts.

Overall, these results are particularly encouraging given the previous discussion highlighting the period right after our chosen estimation window as the one during which the implied predictive VAR might be very mis-specified. Given that even under this scenario we find that the TTDF outperforms the TDF, this suggests that the gains would be even larger for other potential out-of-sample experiments. Furthermore, they confirm that the higher performance of the TTDF does not arise because of excessive risk taking; on the contrary it often results from lower risk-taking in anticipation of bad states of the world.

7 Extensions and Robustness

The full set of results associated with this Section are available in an Online Appendix.

7.1 Heteroskedastic Model for Returns

As discussed in Section 2.4 the assumption of homoskedastic returns could impact our conclusions if $t + 1$ -volatility is expected to be high (low) when the VRP predicts high expected returns at $t + 1$, as we would be over-stating (under-stating) the welfare gains by ignoring this link. We estimated a more general model (equation (6)) and found that the coefficient on the VRP in the variance equation was not statistically significant. Nevertheless, as a robustness exercise, we now explore how the results change if we augment our model for returns with equation (6) and our point estimates of a and b .⁴¹ For the baseline coefficient

⁴¹In our implementation we exclude the error term for computational reasons. The error term would naturally be present in both the VRP and the i.i.d. versions of the model, so it is unlikely to have a meaningful impact on the results.

of risk aversion of 5 we find that the welfare gains are lower, as expected, but the differences are very small. For example, for the case with no additional fee (Extra fee= 0) the age-20 welfare gains are now 5.64%, 2.34%, and 1.32%, respectively, for the three different values of the turnover constraint. These compare with 5.72%, 2.37%, and 1.35% in the baseline case.

7.2 Extended Tactical TDF

The portfolio rule based on equation (26) is straightforward but quite restrictive relative to the optimal model. In this section we 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

$$\alpha_{iat} = I_{f_t=f_j} \cdot \theta_0^j + I_{f_t=f_j} \cdot \theta_1^j \cdot a + \varepsilon_{iat}^j, \text{ for each } f_j \quad (29)$$

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

We obtain corresponding certainty equivalent gains ranging from 5.58% ($\gamma = 5$ and fee=100 basis points) to 12.1% ($\gamma = 10$ and fee=0). By comparison, the corresponding gains for the TTDF were 4.25% and 10.1%. For the reasons that we previously discussed we do not view this rule as a very practical proposition for most TDFs, but these results suggest that individuals with high financial literacy who would potentially be willing to invest in these hypothetical funds could obtain significantly larger CE gains.

7.3 Relaxing the short-selling constraints

We imposed fully binding short-selling constraints on the TTDF because a mutual fund that takes leveraged positions might not be regarded as an acceptable choice by some pension plan providers. We now explore the potential increase in utility gains from the VRP strategy by relaxing those constraints. More precisely 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. 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%.

For a relative risk aversion of 5 and no turnover constraints ($k = 1$), the age-20 welfare gains are now 10.4%, 9.45% and 8.1%, respectively for the different levels of fees. These are substantially larger than the ones reported in Table 6, where short-selling was completely ruled out: 5.72%, 5.12% and 4.25%, respectively. However, if we impose turnover restrictions, such that the level of trading is essentially identical to the one from the baseline TDF, the welfare gains are similar to the ones in Table 6. We conclude that relaxing the short-selling constraint on the riskless asset can increase the welfare gains from the TTDF, but only if we are willing to accept a higher level of turnover.

7.4 Adding VRP strategies during retirement

In our final extension we consider the additional benefits of combining the TTDF with a fund for the retirement period designed in the same manner. More precisely, we run a second regression given by equation (26) 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 welfare gains (both at age 20 and at age 65) are noticeably higher. For example, the age-20 certainty equivalent increases from 5.72% to 6.86% for the unrestricted turnover case with no additional fee and risk aversion of 5.

8 Conclusion

We analyze 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 can

be economically large even after restricting the turnover of the TTDF. These gains could be potentially increased by considering different extensions to the simplified rule or by considering predictive variables with even higher forecasting power, such as the implied correlation or the correlation risk premium. 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.

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] Brennan, M.; E. Schwartz; and R. Lagnado. “Strategic Asset Allocation,” *Journal of Economic Dynamics and Control*, 21 (1997), 1377–1403.
- [10] Brennan, M., and Y. Xia. “Dynamic Asset Allocation under Inflation,” *Journal of Finance*, 57 (2002), 1201–38.
- [11] Buraschi, A., P. Porchia, and F. Trojani, “Correlation risk and optimal portfolio choice,” *The Journal of Finance*, 65, (2010), 393–420.
- [12] Buss, A., L. Schonleber and G. Vilkov, “Expected Correlation and Future Market Returns,” working paper, INSEAD and Frankfurt School of Finance & Management, 2018.
- [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.; J. Cochanre. “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107 (1999), 205–251.
- [17] Campbell, J.; Y. L. Chan; and L. Viceira. “A Multivariate Model of Strategic Asset Allocation,” *Journal of Financial Economics*, 67 (2003), 41–80.
- [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] Dahlquist, Magnus, Offer Setty and Roine Vestman. “On the Asset Allocation of a Default Pension Fund,” *Journal of Finance*, 73, (2018), 1893-1936.
- [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, A., C. Gottlieb and L. Guiso, “Asset Market Participation and Portfolio Choice over the Life-Cycle,” *Journal of Finance*, 72, (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] Flodén, 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] Guiso, L. Sapienza, P. and L. Zingales, “Trusting the Stock Market,” *Journal of Finance* 63 (2008), 2557-2600.
- [30] Johannes, M., A. Korteweg, and N. Polson, “Sequential learning, predictability, and optimal portfolio returns,” *The Journal of Finance* 69 (2014), 611–644.

- [31] Kim, T. S., and E. Omberg. “Dynamic Nonmyopic Portfolio Behavior,” *Review of Financial Studies*, 9 (1996), 141–161.
- [32] 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.
- [33] Lettau, M., and S. Ludvigson, “Consumption, Aggregate Wealth and Expected Stock Returns,” *Journal of Finance*, 56, 3, (2001), 815–849.
- [34] Lettau, Martin, and Stijn Van Nieuwerburgh, “Reconciling the return predictability evidence,” *Review of Financial Studies*, 21 (2008), 1607–1652.
- [35] Liu, J., “Portfolio selection in stochastic environments,” *Review of Financial Studies*, 20 (2007), 1–39.
- [36] Lusardi, A. and O. Mitchell. “The Economic Importance of Financial Literacy: Theory and Evidence,” *Journal of Economic Literature*, 52 (2014), 5-44.
- [37] Lynch, A., and S. Tan. “Labor Income Dynamics at Business-Cycle Frequencies: Implications for Portfolio Choice,” *Journal of Financial Economics*, 101 (2011), 333–359.
- [38] 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.
- [39] 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.
- [40] Moreira, A., and T. Muir, “Volatility-managed portfolios,” *Journal of Finance* 72 (2017), 1611-1644.
- [41] Moreira, A., and T. Muir, “Should Long-Term Investors Time Volatility?” *Journal of Financial Economics* 131-3 (2019), 507-27.
- [42] Munk, C., and C. Sorensen. “Dynamic asset allocation with stochastic income and interest rates,” *Journal of Financial Economics*, 96 (2010), 433–462.

- [43] Pastor, L., and R. F. Stambaugh. “Are Stocks Really Less Volatile in the Long Run?” *The Journal of Finance*, 67 (2012), 431–477.
- [44] 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.
- [45] Tauchen, G., and R. Hussey. “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica*, 59 (1991), 371–396.
- [46] Sialms, C., L. Starks and H. Zhang, “Defined Contribution Pension Plans: Sticky or Discerning Money?”, Working paper, U. Texas at Austin 2013.
- [47] Vissing-Jorgensen, A. “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 110 (2002), 825–853.
- [48] 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.
- [49] Weil, P. “Non-Expected Utility in Macroeconomics,” *Quarterly Journal of Economics*, 2 (1990), 29–42.
- [50] Zhou, H. “Variance Risk Premia, Asset Predictability Puzzles, and Macroeconomic Uncertainty,” *Annual Review of Financial Economics*, (2018), 481-497.

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. CAY is the consumption-asset-labor income variable from Lettau and Ludvigson (2001) for this sample period, and DP is the dividend yield. Inflation (π) is derived from CPI. Inflation, the dividend yield 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	$IV - RV$	CAY	DP	π
Mean (%)	1.98	1.11	0.62	0.49	0.04	0.55	0.60
SD (%)	7.80	0.94	0.98	0.94	2.00	0.16	0.80
Kurtosis	3.24	8.16	54.2	31.8	2.67	3.13	9.64
Skewness	-0.40	2.25	6.45	-3.24	-0.44	0.50	-1.39
AR(1)	0.00	0.41	0.47	-0.17	0.88	0.82	-0.10
Panel B: Correlation Matrix							
1990Q1 –2016Q3	r	IV	RV	$IV - RV$	CAY	DP	π
r	1.00	-0.52	-0.42	-0.10	-0.09	0.23	-0.11
IV	–	1.00	0.70	0.34	0.26	-0.20	-0.18
RV	–	–	1.00	-0.43	0.09	-0.11	-0.46
$IV - RV$	–	–	–	1.00	0.20	-0.11	0.38
CAY	–	–	–	–	1.00	0.20	0.06
DP	–	–	–	–	–	1.00	-0.01
π	–	–	–	–	–	–	1.00

Table 2: Predictive Regressions

Table 2, Panel A, 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 that sets some coefficients equal to zero. The unrestricted equation in the online appendix shows the restricted coefficients are statistically insignificant from zero and we therefore use the restricted VAR below in comparing different models. We follow the same approach when estimating the model with CAY and DP as the predictor variables. Newey-West t-statistics are reported in parentheses.

$$\begin{bmatrix} f_{t+1} \\ r_{t+1} - r_f \end{bmatrix} = \begin{bmatrix} Const \\ \alpha \end{bmatrix} + \begin{bmatrix} \phi & 0 \\ \beta_f & 0 \end{bmatrix} \begin{bmatrix} f_t \\ r_t - r_f \end{bmatrix} + \begin{bmatrix} \varepsilon_{t+1} \\ z_{t+1} \end{bmatrix}$$

Table 2, Panel B presents the estimated relationship between the different factors and the variance of returns from the following restricted VAR that sets some (statistically insignificant) coefficients equal to zero. We also report the results from a monthly frequency VRP for comparison purposes.

$$Var_{t+1} = a + bf_t + v_{t+1}$$

Panel A. Comparing VRP, CAY and DP.

1990Q1 –2016Q3	<i>VRP</i>		<i>CAY</i>		<i>DP</i>	
β_f	3.60	(4.48)	0.55	(1.40)	3.70	(2.83)
ϕ	-0.18	(-1.84)	0.93	(22.7)	0.82	(15.4)
$\rho_{z,\varepsilon}$	-0.04		-0.51		0.22	
σ_ε	0.007		0.008		0.001	
σ_z	0.075		0.078		0.079	
σ_r	0.079		0.079		0.079	
Adj. R^2 (%)	15.0		5.40		6.20	

Panel B. Effect of VRP, CAY and DP on variance on next period returns.

1990Q1 –2016Q3	<i>VRP</i>		<i>VRP (Monthly)</i>	
a	0.0058	(5.03)	0.24	(9.23)
b	0.0840	(0.66)	-0.23	(-2.34)
σ_z	0.0098		0.37	
Adj. R^2 (%)	-0.0055		0.01	
1990Q1 –2016Q3	<i>CAY</i>		<i>DP</i>	
a	0.0062	(6.45)	0.01	(3.80)
b	0.0270	(0.55)	-1.21	(-2.06)
σ_z	0.0098		0.01	
Adj. R^2 (%)	-0.0067		2.90	

Table 3: Household Consumption/Income Growth Sensitivity to Predictors Across Horizons

Table 3 presents the sensitivity of different moments of stockholder quarterly consumption/income growth to the variance risk premium (VRP), dividend yield (DP) and CAY over horizons of $S = 1, 2, 4$ and 8 quarters. Panel A (C) reports the sensitivity of mean consumption (labor income) growth, while Panel B reports the results for the Standard Deviation, Skewness and Kurtosis for consumption growth (not possible to do the same for income using CEX). The consumption growth rate is computed for the households within the 95th percentile of the consumption distribution (details in online appendix). The sensitivity is computed as the regression coefficient from regressing a group's consumption growth over horizon S on the current VRP, DP and CAY. Below each entry we include t-stats. 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 and Different Predictors

Mean Consumption Growth (1996–2015)				
S	1	2	4	8
VRP	1.14	0.63	-0.71	0.79
(t-stat)	(1.60)	(1.14)	(1.57)	(1.41)
DP	-11.3	-23.6	-18.6	-32.4
(t-stat)	(-1.72)	(-3.66*)	(-2.76*)	(-3.43*)
CAY	0.34	0.25	0.28	1.66
(t-stat)	(0.68)	(0.31)	(0.31)	(0.86)

Panel B: Higher Moments of Consumption Growth and Different Predictors

1996–2015	Std. Dev.		Skewness		Kurtosis	
S	1	2	1	2	1	2
VRP	0.43	-0.92	0.24	2.14	-8.30	13.1
(t-stat)	(0.99)	(-1.62)	(0.07)	(0.64)	(-1.09)	(0.93)
DP	-1.79	0.73	-33.1	-28.4	-24.0	-12.7
(t-stat)	(-0.44)	(0.16)	(-1.06)	(-0.89)	(-0.34)	(-0.19)
CAY	0.54	0.39	0.09	-1.23	-9.27	3.99
(t-stat)	(1.80*)	(1.18)	(0.04)	(-0.62)	(-1.77)	(0.76)

Panel C: Mean Income Growth and Different Predictors

Mean Income Growth Rate (1996–2015)				
S	1	2	4	8
VRP	1.50	2.40	0.42	-0.27
(t-stat)	(1.10)	(1.70)	(0.35)	(-0.15)
DP	2.90	-8.5	-12.0	-31.0
(t-stat)	(0.24)	(-0.79)	(-1.20)	(-3.20*)
CAY	-0.026	0.09	0.53	1.29
(t-stat)	(-0.03)	(0.13)	(0.71)	(0.99)

Table 4: Welfare for different values of risk aversion and EIS

Table 4 compares welfare gains and wealth changes when the predictor is either the variance risk premium (VRP), or dividend yield (DP) or CAY relative to the i.i.d. stock returns model. Panel A reports the results for different risk aversion coefficients (2, 5, 10). $W_{65}(\%)$ is the % increase in wealth relative to the i.i.d. model, $CEG_{65}^i(W_{65}^i)$ is the % increase in welfare relative to the i.i.d. model evaluated at VRP/CAY/DP mean wealth at age 65 vs the age 65 i.i.d. model for the same wealth, $CEG_{65}^i(W_{65}^{VRP})$ is the % increase in welfare relative to the i.i.d. model evaluated at VRP mean wealth at age 65 for both models (numerator and denominator), $CEG_{65}^i(W_{65}^i(C_{20-65}^{iid}))$ is the % increase in Age-65 consumption equivalent (CE) gain (explained in the text), CEG_{20}^i is the % increase in lifetime certainty equivalent consumption when comparing each model with the i.i.d. model. Panel B reports the same comparative statics for EIS equal to 1.5 and discount factor equal to 0.995.

Panel A presents model comparisons and consumption certainty equivalent gains when the EIS = 0.5

EIS	0.5								
Discount Factor	0.9875								
Risk Aversion	2			5			10		
Model (i)	VRP	CAY	DP	VRP	CAY	DP	VRP	CAY	DP
$W_{65}(\%)$	161	34.8	19.8	134	50.9	21.1	110	44.1	12.0
$CEG_{65}^i(W_{65}^i)(\%)$	8.83	1.09	0.61	15.9	4.35	0.12	13.7	4.71	0.00
$CEG_{65}^i(W_{65}^{VRP})(\%)$	8.83	2.39	1.25	15.9	6.01	0.26	13.7	5.52	0.10
$CEG_{65}^i(W_{65}^i(C_{20-65}^{iid}))(\%)$	52.3	9.12	0.46	96.9	24.4	6.27	134	45.4	6.33
$CEG_{20}^i(\%)$	2.58	0.59	0.29	8.18	3.22	1.01	12.9	5.40	0.85

Panel B presents the ratio of value function at different incomes when the EIS = 1.5

EIS	1.5			0.5					
Discount Factor	0.9875			0.9875			0.995		
Risk Aversion	5								
Model (i)	VRP	CAY	DP	VRP	CAY	DP	VRP	CAY	DP
$W_{65}(\%)$	502	140	25.1	134	50.9	21.1	112	46.4	18.6
$CEG_{65}^i(W_{65}^i)(\%)$	22.5	4.56	0.50	15.9	4.35	0.12	28.1	9.80	2.36
$CEG_{65}^i(W_{65}^{VRP})(\%)$	22.5	8.00	1.23	15.9	6.01	0.26	28.1	11.2	2.51
$CEG_{65}^i(W_{65}^i(C_{20-65}^{iid}))(\%)$	70.7	12.9	0.00	96.9	24.4	6.27	153	50.5	14.6
$CEG_{20}^i(\%)$	10.4	2.93	0.60	8.18	3.22	1.01	21.8	9.11	2.77

Table 5: Regression for different values of risk aversion and EIS

Table 5 presents the regression of simulated portfolios on age and factor realizations across different relative risk aversion coefficients (5 and 10) for $\psi = 0.5$. More precisely, we use the simulated output from the model to estimate: $\alpha_{iat} = \theta_0 + \theta_1 * a + \theta_2 * f_t + \varepsilon_{iat}$.

$\psi = 0.5$				
	$\gamma = 5$		$\gamma = 10$	
	VRP	IID	VRP	IID
θ_0	0.51 (0.00024)	1.06 (0.0002)	0.26 (0.00022)	0.45 (0.00016)
θ_1	-0.0019 (0.000002)	-0.0031 (0.000002)	-0.0013 (0.000002)	-0.0024 (0.0000016)
θ_2	45.6 (0.015)		43.0 (0.014)	

Table 6: Welfare gains from the TTDF

Table 6 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. In the TTDF the portfolio allocation also depends on the variance risk premium (VRP), which enters as an additional variable in the linear regression. Results are shown for different values of risk aversion (γ), different magnitudes of the additional transaction costs (Extra Fee) faced by the TTDF relative to the TDF, expressed in annualized basis points (0, 40, and 100), and different maximum rebalancing constraints (0.15, 0.25 and 1.0 (no constraints)). $W_{65}(\%)$ is the % increase in mean wealth at age 65, $Sd(W_{65})(\%)$ is the % increase in the cross-sectional standard deviation of wealth at age 65, $CEG_{65}(W_{65}(C_{20-65}^{iid}))(\%)$ is the % increase in Age-65 consumption equivalent (CE) gain (using the i.i.d. consumption rule, more details in the text), and CEG_{20} is the % increase in lifetime certainty equivalent consumption.

Risk Aversion	5								
Max Rebalancing	1.0	1.0	1.0	0.25	0.25	0.25	0.15	0.15	0.15
Extra Fee	0	40	100	0	40	100	0	40	100
Mean Turnover	209	209	209	103	103	103	66.2	66.2	66.2
$W_{65}(\%)$	183	161	131	40.0	27.7	10.8	8.99	-1.14	-15.0
$Sd(W_{65})(\%)$	247	222	187	40.2	26.9	9.49	-0.93	-10.5	-22.6
$CEG_{65}(W_{65}(C_{20-65}^{iid}))(\%)$	40.2	34.1	26.0	9.80	6.06	1.04	2.41	-0.16	-5.03
$CEG_{20}(\%)$	5.72	5.12	4.25	2.37	1.85	1.20	1.35	0.89	0.25
Risk Aversion	10								
Max Rebalancing	1.0	1.0	1.0	0.25	0.25	0.25	0.15	0.15	0.15
Extra Fee	0	40	100	0	40	100	0	40	100
Mean Turnover	181	181	181	97.4	97.4	97.4	62.8	62.8	62.8
$W_{65}(\%)$	309	291	263	137	120	103	85.0	74.9	60.5
$Sd(W_{65})(\%)$	497	482	458	255	236	209	169	153	130
$CEG_{65}(W_{65}(C_{20-65}^{iid}))(\%)$	78.3	71.7	62.5	35.8	31.4	25.4	23.5	19.9	14.8
$CEG_{20}(\%)$	10.1	9.23	7.95	4.84	4.03	2.85	2.84	2.06	0.92

Table 7: Welfare gains from the TTDF with Parameter Uncertainty

Table 7 presents results from comparing the TTDF with the standard TDF, when the TTDF rules incorporate parameter uncertainty as described in the main text. In both models the i.i.d. consumption rule is used. The results are for the investor with risk aversion equal to 5, with different rebalancing restrictions (0.15, 0.25 and 1.0 (no constraints)), and for different magnitudes of the additional transaction costs faced by the TTDF relative to the TDF (Extra Fee), expressed in annualized basis points. $W_{65}(\%)$ is the % increase in mean wealth at age 65, $Sd(W_{65})(\%)$ is the % increase in the cross-sectional standard deviation of wealth at age 65, $CEG_{65}(W_{65}(C_{20-65}^{iid}))$ is the Age-65 consumption equivalent (CE) Gain (using the i.i.d. consumption rule, more details in the text), and CEG_{20} is the % increase in lifetime certainty equivalent consumption.

Maximum Rebalancing	1.0			0.25			0.15		
Extra Fee	0	40	100	0	40	100	0	40	100
Mean Turnover	207	207	207	103	103	103	66.2	66.2	66.2
$W_{65}(\%)$	181	159	129	40.0	27.7	10.8	9.07	-1.07	-14.9
$Sd(W_{65})(\%)$	246	221	184	40.4	27.1	9.55	-0.75	-10.2	-22.5
$CEG_{65}(W_{65}(C_{20-65}^{iid}))(\%)$	40.2	34.1	24.2	9.79	6.05	-0.08	2.40	-0.75	-6.06
$CEG_{20}(\%)$	5.64	5.08	4.09	2.36	1.85	1.01	1.31	0.81	0.02

Table 8: Predictive Regressions for Different Sub-Periods

Table 8 presents a restricted VAR based on quarterly data for different sub-periods of the sample: 1990:1 to 1999:4, 2000:1 to 2009:4 and 2010:1 to 2016:3. The restricted VAR is given by:

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

Newey-West t-statistics are reported in parentheses.

	1990Q1–1999Q4	2000Q1–2009Q4	2010Q1–2016Q3
β_f	5.85 (4.88)	2.01 (1.62)	5.23 (2.85)
ϕ	0.18 (1.15)	-0.30 (-2.01)	-0.12 (-0.60)
$\rho_{z,\varepsilon}$	-0.41	0.12	-0.61
σ_ε	0.005	0.010	0.006
σ_z	0.064	0.089	0.065
σ_r	0.070	0.092	0.071
Adj. R^2 (%)	34.6	3.60	21.4

Figure 1 shows the time series of implied volatility (IV), realized volatility (RV) and variance risk premium (VRP). The series for realized volatility is taken from Zhou (2017) and based on daily US stock market returns from CRSP, while the series for implied volatility is taken from the Federal Reserve Bank of St. Louis. The variance risk premium is the difference between the other two. All data are quarterly from 1990 to 2016.

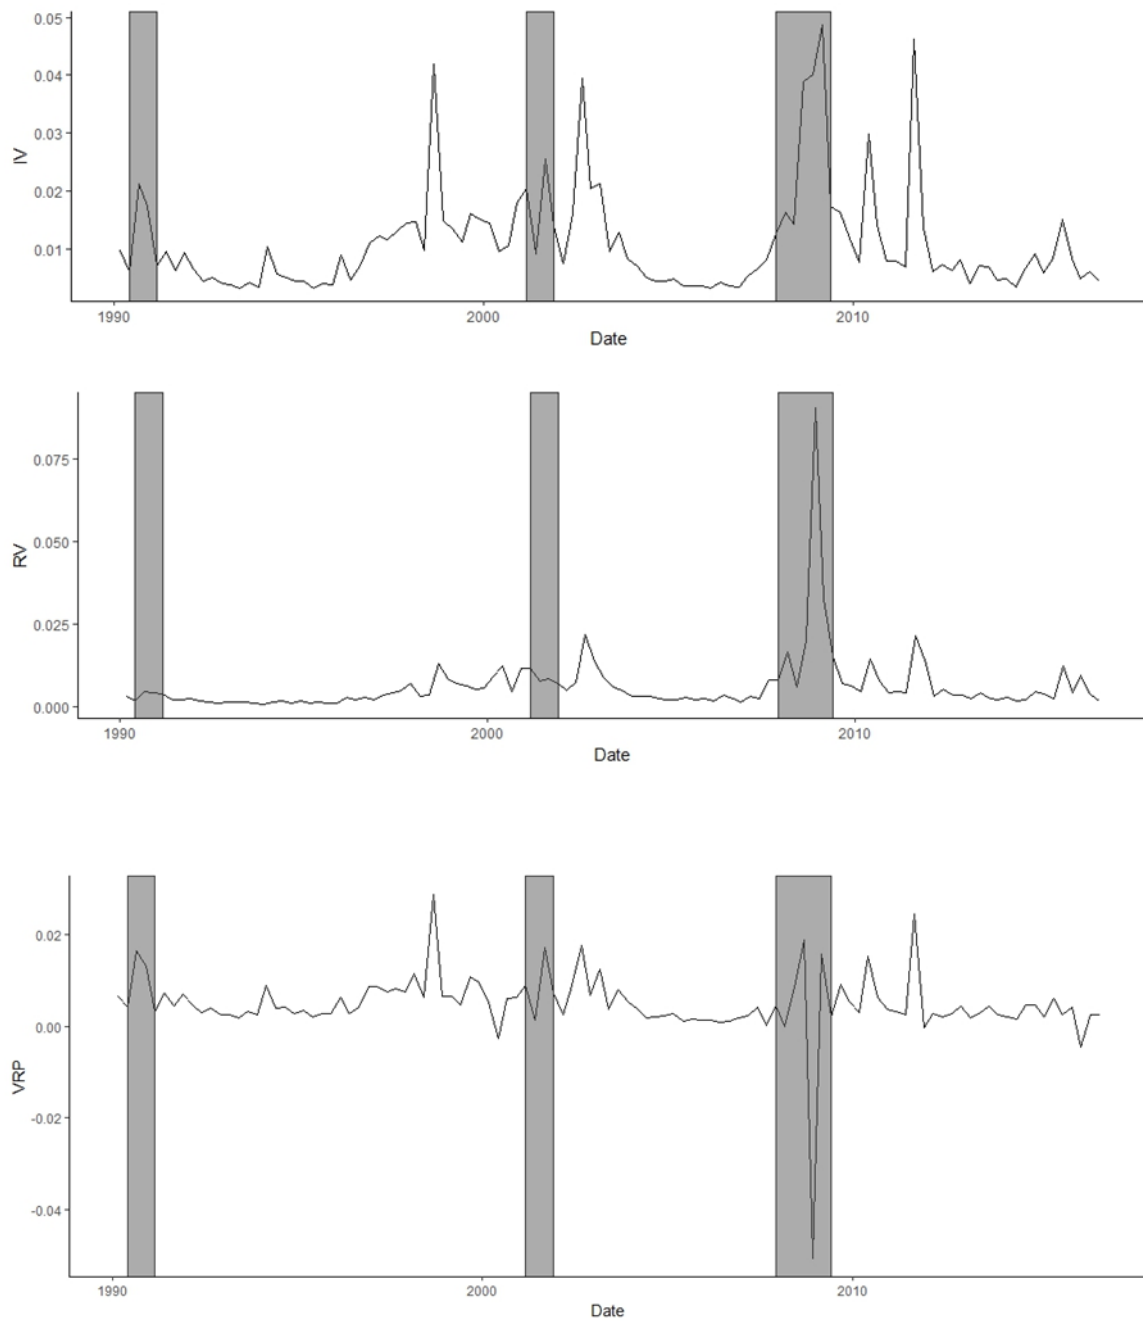


Figure 2.1 shows the optimal pre-retirement portfolio allocations both for the investor using the i.i.d. model for returns (“i.i.d. investor”) and for the investor using the VAR model for returns with the VRP predictor (“VRP investor”). For the VRP investor we report both the average allocation and the allocation for the average realization of the predictive factor.




Figure 2.2 shows the mean optimal allocation for the VRP investor and the allocation at plus/minus one standard deviation from the mean.

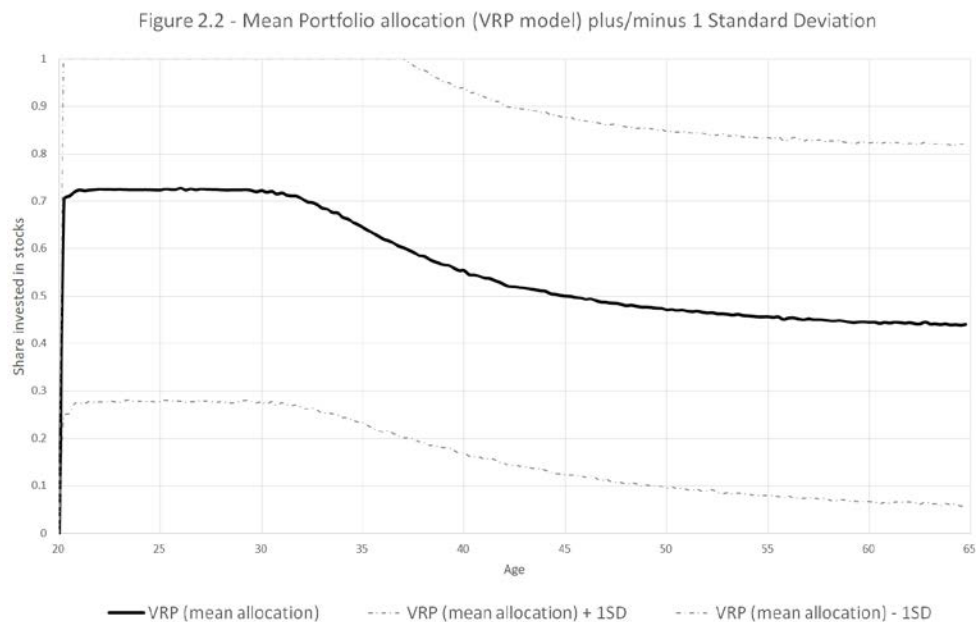


Figure 3.1 shows the expected portfolio return both for the investor using the i.i.d. model for returns (“i.i.d. investor”) and for the investor using the VAR model with the VRP predictor (“VRP investor”). To facilitate a comparison, we also plot the difference between the two.




Figure 3.2 shows the expected portfolio returns both for the investor using the i.i.d. model for returns (“i.i.d. investor”) and for the investor using the VAR model with the VRP predictor (“VRP investor”), and the portfolio returns at plus/minus one standard deviation from the mean.




Figure 4 shows the mean life-cycle portfolio allocation of both the Target-Date Fund (TDF) and the Tactical Target-Date Fund (TTDF). For comparison we also report the average optimal asset allocation of the investor using the i.i.d. model for returns (“i.i.d. investor”) and for the investor using the VAR model for returns with the VRP predictor (“VRP investor”).




Figure 5 shows the life-cycle portfolio allocation of the Tactical Target-Date Fund (TTDF), both with and without turnover restrictions for different realizations of the predictive factor (f). The value of $f=0.494\%$ corresponds to its unconditional mean. The values of 1.35% and -0.36% correspond to 1.14 standard deviations above and below the mean, respectively. The case with turnover restrictions considers a maximum turnover limit (k) of 0.15 per quarter.

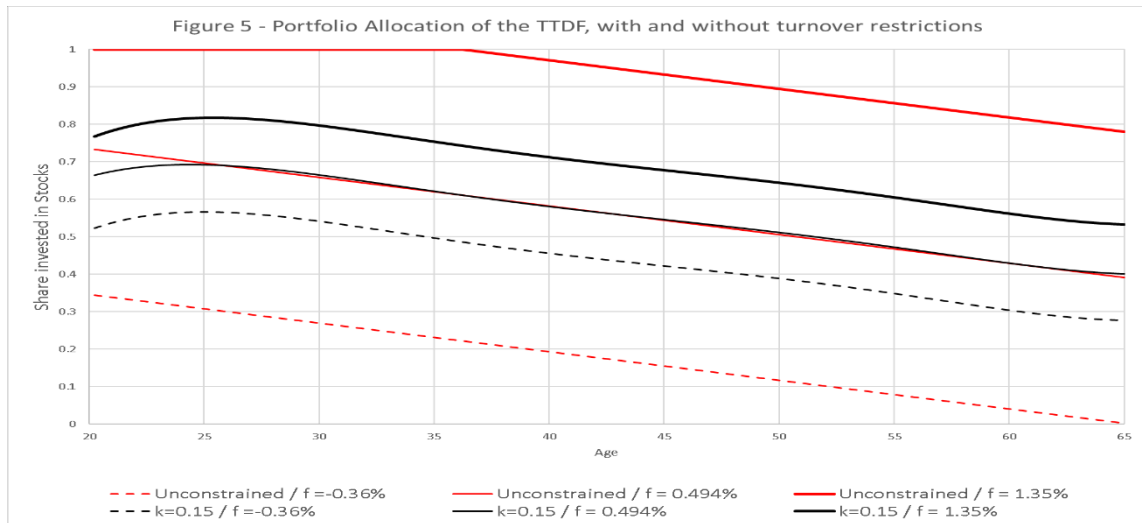


Figure 6.1 shows the results of an out-of-sample comparison between the TTDF (with a 15% turnover restriction) and the TDF. The figure reports the cumulative return to the TTDF and TDF funds for the 20-year investor from 2000 onwards, when the portfolio allocation of TTDF fund is based on an estimation of the predictive model that only uses data until 1999.




Figure 6.2 shows the results of an out-of-sample comparison between the TTDF (with a 15% turnover restriction) and the TDF. The figure reports the annualized period-by-period return to the TTDF and TDF funds for the 20-year investor from 2000 onwards, when the portfolio allocation of TTDF fund is based on an estimation of the predictive model that only uses data until 1999.




Figure 7.1 shows the results of an out-of-sample comparison between the TTDF (with a 15% turnover restriction) and the TDF. The figure reports the cumulative return to the TTDF and TDF funds for the 50-year investor from 2000 onwards, when the portfolio allocation of TTDF fund is based on an estimation of the predictive model that only uses data until 1999.




Figure 7.2 shows the results of an out-of-sample comparison between the TTDF (with a 15% turnover restriction) and the TDF. The figure reports the annualized period-by-period return to the TTDF and TDF funds for the 20-year investor from 2000 onwards, when the portfolio allocation of TTDF fund is based on an estimation of the predictive model that only uses data until 1999.

