Stock Market Mean Reversion and Portfolio Choice over the Life Cycle*

Alexander Michaelides[†]

Imperial College London, CEPR and NETSPAR

Yuxin Zhang

Imperial College London[‡]

November 6, 2015

*This version is preceded by an older working paper entitled "Portfolio Choice, Liquidity Constraints and Stock Market Mean Reversion." We would like to thank an anonymous referee, Nicholas Barberis, John Campbell, John Cochrane, George Constantinides, Steve Davis, Christian Gollier, Francisco Gomes, Michael Haliassos, James Kahn, Martin Lettau, Sydney Ludvigson, Lubos Pastor, Cesare Robotti, Annette Vissing-Jorgensen, Paul Willen and seminar participants at the Athens University of Economics and Business, the NETSPAR international conference on pensions, the University of Chicago Graduate School of Business, the University of Cyprus, the University of Southampton, the University of Sussex, and the New York Federal Reserve Bank for many helpful comments.

[†]Address: Department of Finance, Imperial College Business School, South Kensington Campus, SW7 2AZ: a.michaelides@imperial.ac.uk.

[‡]Address: Department of Finance, Imperial College Business School, South Kensington Campus, SW7 2AZ: y.zhang11@imperial.ac.uk.

Stock Market Mean Reversion and Portfolio Choice over the Life Cycle

Abstract

We solve for optimal consumption and portfolio choice in a life-cycle model with short-sales and borrowing constraints, undiversifiable labor income risk and a predictable, time-varying, equity premium and show that the investor pursues aggressive market timing strategies. Importantly, in the presence of stock market predictability, the model suggests that the conventional financial advice of reducing stock market exposure as retirement approaches is correct on average, but ignoring changing market information can lead to substantial welfare losses. Therefore, enhanced target-date funds (ETDFs) that condition on expected equity premia increase welfare relative to target-date funds (TDFs). Out-of-sample analysis supports these conclusions.

JEL Classification: E21, G11.

Key Words: Portfolio Choice over the Life Cycle, Stock Market Mean Reversion, Stock Market Predictability, Hedging Demands, Lifestyle funds, enhanced target-date funds.

I Introduction

In the absence of stock market predictability and in the presence of borrowing and short sales constraints,¹ models with background labor income risk (e.g., Cocco, Gomes, and Maenhout (2005)) predict that households should invest a larger proportion of their savings in the stock market when young: future labor income acts like an implicit risk free asset crowding out riskless financial assets from household portfolios. This prediction resembles the advice given by financial planning consultants in recommending target-date funds (TDFs) that reduce exposure to stocks as retirement approaches.

How does the presence of stock market predictability affect these predictions? Various papers have analyzed the implications of stock market predictability for consumption and/or portfolio choice while ignoring labor income risk,² whereas others focus on the effect of background labor income risk on portfolio choice while ignoring stock market predictability.³ This paper jointly models stock market predictability and non-diversifiable background labor income risk and analyzes the normative implications for optimal consumption and portfolio choice over the life cycle using Epstein-Zin (1989) and Weil (1990) preferences.⁴ The results

¹Deaton (1991) and Carroll (1997) propose models with background risks and potentially binding liquidity constraints (“buffer stock saving models”) as the leading alternative to the classic Permanent Income-Life Cycle Hypothesis.

²Kim and Omberg (1996), Brennan, Schwartz and Lagnado (1997), Brandt (1999), Campbell and Viceira (1999), Balduzzi and Lynch (1999), Barberis (2000), Campbell et. al. (2001), (2003), and Wachter (2002) show that stock market exposure varies substantially as a response to the predictive factor(s).

³For instance, Heaton and Lucas (2000), Viceira (2001), Haliassos and Michaelides (2003), Gomes and Michaelides (2005), and Cocco, Gomes, and Maenhout (2005).

⁴Benzoni, Collin-Duffresne, and Goldstein (2007) investigate the implications of a cointegrating relationship between labor income and stock returns to show how stock demand for young investors can be reduced relative to the absence of this type of long-run risk. Lynch

make the case for enhanced target-date funds (ETDFs) that condition decisions on expected risk premia: introducing such funds generates improvements in investor welfare.

As in the buffer stock saving literature, optimal consumption is a concave function of liquid wealth, but does not respond substantially to changes in the investment opportunity set. On the other hand, the consumer/investor is shown to be an aggressive market timer in the presence of stock market predictability. Relative to the i.i.d. returns model, high expected future returns generate a higher allocation of stocks in the portfolio for a given level of saving (when constraints are not binding), while low expected future returns decrease the exposure in the stock market. This translates to large variations over the life cycle depending on the factor realization rather than the level of financial wealth. This result substantially alters one of the main insights of life-cycle models with i.i.d. stock returns, namely that financial wealth tends to be the main predictor of life-cycle portfolio choice.

When calibrated to the observed dividend yield as a factor of predictability, the asset allocation profiles retain the target date funds (TDF) feature of slowly decreasing stock market exposure as the household ages. Nevertheless, the level of asset allocation moves up or down depending on the factor realization: optimal portfolio choice shifts up or down depending on the expected risk premium. When experimenting with a more volatile process, the portfolio movements become a lot more aggressive to take advantage of factor predictability.⁵ Taken and Tan (2011) generate a similar result by focussing on the implications of time-variation in the mean and, in particular, the variance of labor income. Munk and Sorensen (2010) focus instead on time variation in interest rates and expected income growth to illustrate the effects on portfolio choice, while Koijen, Nijman, and Werker (2010) focus on the effects of bond risk premia predictability on optimal life cycle asset allocation. Brennan and Xia (2002) instead focus on the effects of inflation on dynamic asset allocation.

⁵Aggressive market timing behavior is similar to the behavior predicted in Brennan,

together, these findings make the case for enhanced TDFs (ETDFs) that condition on the market timing ability of the investor.

From all the underlying correlations studied, the main correlation that is found to quantitatively affect our conclusions is the correlation between permanent earnings shocks and the stock market innovation. Financial advisors should therefore pay special attention to this correlation when devising rules of thumb about life cycle portfolio allocations. We also experiment with introducing model uncertainty with regards to the persistence parameter and the correlation between the factor and stock market innovation and find that they do not substantially alter the asset allocation and wealth accumulation profiles.

To emphasize the results with regards to ETDFs, we next make welfare comparisons across the Vanguard TDF recommendations, an i.i.d. stock returns model and the baseline mean reversion model. We show substantial welfare losses relative to the baseline and we show how these can arise either from lower mean consumption or higher consumption volatility over the life cycle when the incorrect portfolio rule is being used. Moreover, these losses are maximized at around age 50 when the increase in average wealth accumulation slows down and the net saving rate (the difference between labor income and consumption) turns negative. We therefore experiment with rules of thumb that alter the Vanguard recommendation depending on the expected equity premium and show how these rules of thumb (a proxy for ETDFs) reduce welfare less than following either the i.i.d. model or Vanguard's recommendation.

We also perform out-of-sample analysis adapting Lan's (2015) methodology to our life Schwartz, and Lagnado (1997) and Barberis (2000), models that do not feature undiversifiable labor income uncertainty.

cycle setting and using the realized dividend yield and stock returns to simulate life cycle wealth and utility in two subperiods: 1985 to 2014 and 1999 to 2014. We find that welfare is statistically and economically significantly higher when comparing the baseline model to either the Vanguard or i.i.d. stock returns model, especially in the 1999 to 2014 period. The analysis demonstrates the value in devising ETDFs for individual investors and calls for further research in this area.

The paper is organized as follows. Section II describes the theoretical model, outlines the numerical solution algorithm and discusses the parameter choices for the calibration. Section III discusses the effects of stock market mean reversion by comparing the benchmark results to the i.i.d. stock returns model. Section IV discusses hedging demands and how different correlation changes also affect wealth accumulation, while Section V briefly discusses the implications of model parameter uncertainty. Section VI discusses the implications of the model for lifestyle funds and Section VII concludes.

II The Model

Time is discrete, there is one non-durable good, one riskless financial asset and a risky time varying investment opportunity. The riskless asset yields a constant gross after tax real return, R_f , while the gross real return on the risky asset is denoted by $\tilde{R}$. At time t , the agent enters the period with invested wealth in the stock market S_{t-1} and the bond market B_{t-1} and receives Y_t units of the non-durable good. Following Deaton (1991), cash on hand in period t is denoted by $X_t = S_{t-1}\tilde{R}_t + B_{t-1}R_f + Y_t$. The investor then chooses savings in the bond (B_t) and stock (S_t) market to maximize welfare. The particular assumptions made

about the economic environment are as follows:

A Preferences

Preferences separate the elasticity of intertemporal substitution from risk aversion as in Epstein and Zin (1989) and Weil (1990). Specifically, they are given by

$$(1) \quad V_t = \left\{ (1 - \beta)C_t^{1-1/\psi} + \beta \left(E_t(p_{t+1}V_{t+1}^{1-\gamma} + b(1 - p_{t+1})X_{t+1}^{1-\gamma}) \right)^{\frac{1-1/\psi}{1-\gamma}} \right\}^{\frac{1}{1-1/\psi}},$$

where β is the time discount factor, b is the strength of the bequest motive, ψ is the elasticity of intertemporal substitution (EIS) and γ is the coefficient of relative risk aversion. The conditional probability of surviving next period conditional on having survived until period t is given by p_{t+1} .

B Labor Income Process

Following a relatively standard specification in the literature (Carroll (1997)), the labor income process before retirement is given by

$$(2) \quad Y_{it} = Y_{it}^p U_{it},$$

$$(3) \quad Y_{it}^p = \exp(g(t, Z_{it}))Y_{it-1}^p N_{it},$$

where $g(t, Z_{it})$ is a deterministic function of age and household characteristics Z_{it} , Y_{it}^p is a permanent component with innovation N_{it} , and U_{it} a transitory component of labor income, where $\ln U_{it}$ and $\ln N_{it}$ 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. The log of Y_{it}^p evolves as a random walk with a deterministic drift, $g(t, Z_{it})$. For simplicity, retirement is assumed to be exogenous and deterministic, with all households retiring in time period K , corresponding to age 65 ($K = 46$). Earnings in retirement ($t > K$) are given by $Y_{it} = \lambda Y_{iK}^p$, where λ is the replacement ratio ($\lambda = 0.68$) of the last working period permanent component of labor income.

Durable goods, and in particular housing, can provide an incentive for higher spending early in life. We exogenously subtract a fraction of labor income every year allocated to durables (housing), and this fraction includes both rental and mortgage expenditures. This empirical process is taken from Gomes and Michaelides (2005) and is based on Panel Study Income Dynamics (PSID) data. We choose not to model explicitly the returns from housing following the empirical evidence (e.g., Cocco and Lopes (2015) and references therein) that households tend not to decumulate housing as fast as life-cycle models predict. A prominent explanation tends to be a psychological benefit from continuing to own one's house, an explanation that is consistent with the low observed demand for home equity conversion mortgages (Davidoff (2015)). For these reasons we do not explicitly model the potential effects of housing returns, and focus instead only on investments of liquid financial wealth for rich households (that empirically tend to be both stockholders and homeowners).

C Liquidity Constraints

Borrowing and short sales of stocks are not allowed: $B_t \geq 0$ and $S_t \geq 0$ to avoid the counterfactual implication that households lever up to invest in the stock market. The share of wealth in stocks (α_t) is defined as $S_t/(S_t + B_t)$.

D Mean Reversion

We follow Campbell and Viceira (1999) and Pastor and Stambaugh (2012) in assuming that there is a single factor that can predict future excess returns. Letting $\{r_f, r_t\}$ denote the net risk free rate and the net stock market return, respectively and f_t the factor that predicts future excess returns, we have

$$(4) \quad r_{t+1} - r_f = f_t + z_{t+1},$$

$$(5) \quad f_{t+1} = \mu + \phi(f_t - \mu) + \varepsilon_{t+1},$$

where the two innovations $\{z_{t+1}, \varepsilon_{t+1}\}$ are i.i.d. Normal variables with mean equal to zero and variances σ_z^2 and σ_ε^2 , respectively. Contemporaneous correlation between these innovations is allowed, while correlation between the permanent earnings innovation ($\ln N_t$) and $\{z_t, \varepsilon_t\}$ can also exist. Mean reversion in the stock market is captured by the autoregressive nature of the factor (f_t) predicting stock market returns ($\phi > 0$) and negative correlation between the excess stock market return innovation (z_{t+1}) and the innovation to the factor (ε_{t+1}). One of the key contributions of the paper is to understand how changing these correlations affects

saving and portfolio choice decisions over the life cycle. We will also be reporting results from a model with i.i.d. excess returns; in that case $r_{t+1} - r_f = \mu + z_{t+1}$. 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.⁶

E Numerical Solution

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_{it}^p). Letting lower case letters denote variables normalized by the permanent component of labor income (Y_{it}^p), the evolution of the single endogenous state variable is then given by

$$(6) \quad x_{it+1} = \frac{Y_{it}^p}{Y_{it+1}^p} (r_{t+1} \alpha_{it} + r_f (1 - \alpha_{it})) x_{it} + U_{it+1}.$$

An Online Appendix (available at www.jfqa.org) details the numerical solution technique and the numerical accuracy in the implementation of the Tauchen (1986) and Tauchen and Hussey (1991) approximation procedure for a vector autoregression. Numerically, this proves to be a substantial challenge because of the strong persistence in the factor f_t that requires a substantial number of grid points to replicate key moments: the persistence, the conditional variance and the conditional (one-period ahead) equity risk premium. Moreover, the conditional expected risk premium must be positive, especially when the factor is denoting the dividend yield, a theoretical restriction that is not typically imposed in most estima-

⁶The i.i.d. returns specification is the one found in recent papers with either CRRA preferences (Cocco, Gomes, and Maenhout (2005)) or Epstein-Zin-Weil preferences (Gomes and Michaelides (2005) or Cooper and Zhu (2014)).

tion methods (a recent counter example is Pettenuzzo, Timmermann, and Valkanov (2014)). Given the low estimated volatility of the factor, we impose the theoretical restriction that the conditional expected equity premium always be positive in the baseline model. To do so, we follow Floden (2008) and use Tauchen’s (1986) approximation for the very persistent AR(1) case. This allows us to explicitly control the range of the factor (thereby ensuring it is always positive) irrespective of the number of grid points chosen to replicate the chosen persistence and cross correlations. For the case with higher volatility that allows negative factor states we revert to the Tauchen and Hussey (1991) approximation. The online appendix provides extensive details from different experiments to test numerical accuracy.

F Parameter Choice

Even though empirical predictability studies are typically done at a monthly or quarterly frequency, we solve the model at an annual frequency to maintain comparability with the life-cycle literature. The net constant real interest rate, r_f , equals 0.02. Carroll (1997) estimates the variances of the idiosyncratic shocks using data from the *Panel Study of Income Dynamics*, and the benchmark simulations use values close to those: 0.1 for σ_u and 0.1 for σ_n . The deterministic component of labor income is identical to the one used by most life cycle papers in this literature (Cocco, Gomes, and Maenhout (2005)), and this also facilitates comparisons between this model and its counterpart with i.i.d. stock returns. The relatively large estimate for the replacement ratio during retirement (68% of last working period labor income) arises from using both social security and private pension accounts to estimate the benefits in the PSID data and is consistent with not explicitly modelling tax-deferred saving through retirement accounts.

The baseline preference specification is taken to capture the observed behavior of stockholders. Gomes and Michaelides (2005) argue that this is well achieved when using a coefficient of relative risk aversion (γ) equal to 5, and an elasticity of intertemporal substitution (ψ) equal to 0.5. These choices are consistent with the empirical estimates for the elasticity of intertemporal substitution in Vissing-Jorgensen (2002) and the empirical preference parameter estimates in Gomes, Michaelides, and Polkovnichenko (2009). The bequest parameter is set to 2.5 to capture the empirical observation that few rich stockholders die with zero financial assets. We set the discount factor (β) equal to 0.96.

To calibrate the stock market predictability parameters we use one of the more popular predictors of stock returns, namely the dividend yield. We estimate the mean reversion system (4) and (5) using three different data sources: the data from Pastor and Stambaugh (2012),⁷ the publicly available data set on Robert Shiller’s website,⁸ and by decomposing the total return CRSP value weighted index into the capital gain and dividend return part. All of these data series are annual, as is the model frequency, and start at 1802, 1871 and 1926, respectively. All three estimations of the model generate a very persistent factor predicting the real log return on the U.S. equity market. This is relatively stable across different subperiods and we therefore set this persistence parameter at $\phi = 0.9$. The unconditional stock market volatility is given by the unconditional standard deviation of stock returns and is set equal to 0.18.

A key parameter turns out to be the correlation between the factor and the return innovation ($\rho_{z,\varepsilon}$). For the Siegel and CRSP data sets this parameter is estimated at -0.6 ,

⁷We thank Lubos Pastor for kindly providing this data set.

⁸At <http://www.econ.yale.edu/~shiller/data.htm>.

while for the Shiller data set, we estimate it at around zero. Most estimates in the literature are towards the higher negative number (Campbell and Viceira (1999) for a quarterly estimation and Pastor and Stambaugh (2012) for both an annual and quarterly estimation). We therefore use $\rho_{z,\varepsilon} = -0.6$ for the baseline model but also experiment with the $\rho_{z,\varepsilon} = 0.0$ estimate, that also motivates a case with uncertainty about this correlation in what follows.

The factor innovation is very smooth and we estimate (and use) $\sigma_\varepsilon = 0.007^9$ for the baseline model but we also find it useful to compare results with the case when the factor is more volatile ($\sigma_\varepsilon = 0.015$). This second case is interesting because it allows us to admit the possibility of negative expected returns (the factor can be negative), which could be a possibility for a small investor taking prices as given (and using a factor other than the dividend yield). The other benchmark parameters for the generation of stock market returns are $\mu = 0.04$ (an unconditional equity premium at 4%) as in most of the life-cycle literature. Given these estimates, we can infer that the unconditional variance of the stock market return innovation equals $\sigma_z^2 = 0.18^2 - \sigma_f^2$.

It should be noted that no estimate of the correlation between the innovation in the factor predicting stock returns and permanent, idiosyncratic earnings shocks ($\rho_{n,\varepsilon}$) exists in the literature and we therefore set this correlation equal to zero. 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

⁹When discretizing the factor with 15 states using the Tauchen (1986) method with $\sigma_\varepsilon = 0.007$, the actual factor states are 0.0079, 0.0125, 0.0171, 0.0216, 0.0262, 0.0308, 0.0354, 0.0400, 0.0446, 0.0492, 0.0538, 0.0584, 0.0629, 0.0675, 0.0721.

permanent labor income shocks and stock returns ($\rho_{n,z}$) is set equal to 0.15, consistent with the mean estimates in most empirical work (Campbell et. al. (2001) or Davis, Kubler, and Willen (2006)). Nevertheless, this can vary and be higher across heterogeneous occupations (Angerer and Lam (2009)) and/or workers (Bonaparte, Korniotis, and Kumar (2014)) and we therefore experiment with values up to 0.5.

III The Effects of Stock Market Mean Reversion

How does the presence of a factor predicting returns affect saving and portfolio choice behavior relative to the i.i.d. model?

The online appendix presents the consumption function that has the familiar shape from the buffer stock saving literature without risky asset choice (Deaton (1991) and Carroll (1997)): below a cutoff point x^* no saving takes place, while the marginal propensity to consume falls quickly beyond x^* . Once the constraint stops binding, the saving reaction to the factor realization is not quantitatively important.

The policy functions for the share of wealth in stocks illustrate the large dependence of portfolios on the factor realization (Figure 1 for age 25 (Graph A), age 55 (Graph B) and age 75 (Graph C)). Portfolios can shift from 40% to 100% and vice versa depending on the factor realization and the household's age, something that does not happen in the i.i.d. model. The median factor state resembles closely the i.i.d. case, while higher factor realizations (higher expected stock returns) shift the share of wealth in stocks upwards. These policy function graphs illustrate clearly that portfolio allocations should be a lot more volatile in the mean reversion than in the i.i.d. model, while the range of outcomes is wider when households

have longer horizons (comparing age 25 with age 55 or age 75, for example). Moreover, the presence of labor income prevents the portfolio from depicting the yo-yo type behavior found in Brennan, Schwartz, and Lagnado (1997) where the portfolios move from zero to one depending on the factor realization. This happens in the baseline case because the factor realization is always constrained to be positive, a feature that arises from the low volatility of the dividend yield.

Figure 2 illustrates the differences in simulated profiles between the i.i.d. and the mean reversion model, while also introducing the life style recommendations issued by Vanguard (Donaldson et. al. (2013)) for target date funds (TDFs). We introduce Vanguard in anticipation of the discussion that follows on the optimality of TDF advice. Vanguard’s basic recommendation is to invest 90% of a household’s financial wealth in equities until age 40, and start decreasing that share as retirement approaches to reach 50% at age 65. To simulate wealth profiles for this case, we take the portfolio rule as exogenous but the household still makes optimal consumption-saving decisions, taking this portfolio decision as given.

Average wealth accumulation¹⁰ is similar across all three models (Graph A of Figure 2). This is slightly surprising since Graph B illustrates that the average share of wealth in stocks can differ substantially across models, even though average behavior still follows the life style fund intuition: reduce exposure to the stock market as retirement approaches. The fact that the average share of wealth in stocks is never one might be surprising given the results in the i.i.d. version of the model. This arises here because we are simulating based on different initial realizations of the factor and then averaging over them. For most of these factor states, the investor deviates from the 100% asset allocation to stocks.

¹⁰The model is simulated 2000 times for 500 individual life histories.

The deviation from the i.i.d. model is even larger when the factor is perceived to be more volatile ($\sigma_\varepsilon = 0.015$). This is illustrated in Graph D of Figure 2: the average share of wealth in stocks is now substantially below one over the largest part of the life cycle. This reflects the fact that the household now expects negative expected risk premia for the lowest factor realizations, generating a full allocation in cash over this range.

The effect of the factor on portfolios can be more clearly seen by tracking individual portfolio simulations starting from different initial factor realizations. Figure 3 plots what happens over the lifecycle when starting from the lowest (state = 1), sixth, tenth and highest (state = 15) factor realization for the baseline case. Because the factor is persistent, it takes a substantial amount of time for a change to happen: when it does happen, the portfolio moves relatively quickly. Nevertheless, the change is different from the bang-bang movement in the share of wealth in stocks found in Brennan, Schwartz, and Lagnado (1997) that features no labor income risk since the zero allocation in cash is never visited in these simulations. This happens in the baseline model because of the low volatility of the factor and the fact that the factor is always constrained to be positive. In the model with higher factor variability, the bang-bang portfolio behavior between zero to full allocation in the stock market arises once again.

We have also experimented with increasing the volatility of the stock return innovation (σ_z) by 10%, 20% and 50% as a way to capture the uncertainty over the coefficient relating the equity premium to the dividend yield. We find that the share of wealth in stocks is reduced throughout the life cycle relative to the baseline mean reversion case but portfolio behavior remains qualitatively unchanged and do not report these results.

IV Hedging Demands

How do these results change when the correlations between the different innovations vary?

We next use the model to quantitatively assess the magnitude of such hedging demands.

A Variation in Correlations

To investigate the importance of hedging demand due to $\rho_{z,\varepsilon}$ (the correlation between the factor and the stock market innovation), we set it equal to zero. In the Shiller dividend yield data set, we have found this correlation to be close to zero, contrary to our findings using CRSP data. Setting the correlation to zero can therefore be useful in assessing the possible range of hedging demands that might be generated from this correlation.

We also evaluate hedging demands when changing the correlation between permanent earnings shocks and stock market innovations ($\rho_{z,n}$). In our baseline model we use 0.15 for this correlation, a value that reflects the substantial idiosyncratic risk that exists in labor income data. Nevertheless, one cannot deny that there are some households for whom this correlation is substantially higher. Bonaparte, Korniotis, and Kumar (2014) find that this correlation can vary for different households from -0.6 to 0.6 and we therefore use 0.5 to investigate how the results change.

There is no known empirical estimate for $\rho_{n,\varepsilon}$ (the correlation between the factor innovation and the permanent labor income shock) in the literature. There are potentially some a priori reasons to expect it not to be statistically different from zero since earnings shocks at the household level have a large idiosyncratic variance component, yet the component of this variance that can be attributed to aggregate shocks is generally quite small. Given

that most factors predicting stock returns have very low volatility, we expect the correlation between the factor innovation and the much more volatile idiosyncratic earnings shocks to be close to zero. Nevertheless, we consider the potential effects of this correlation by increasing it to a high enough value that can simultaneously maintain the positive definiteness of the variance covariance matrix of the different innovations: setting $\rho_{n,\varepsilon}$ equal to 0.15 satisfies this constraint.

The wealth accumulation and mean shares of wealth in stocks over the life cycle are depicted in Graph A of Figure 4, and Graph B of Figure 4, respectively.¹¹ When the correlation between the factor and the stock return innovation is set to zero ($\rho_{z,\varepsilon} = 0.0$) from $\rho_{z,\varepsilon} = -0.6$, the share of wealth in stocks is lower on average between 0 and 5% over different parts of the working life cycle (Graph B) and the effects are larger when the factor is more volatile (Graph D), leading to a discernibly lower average wealth accumulation (Graph C). It can be shown that the variance of longer term stock returns is higher when $\rho_{z,\varepsilon} = 0.0$ rather than when $\rho_{z,\varepsilon} = -0.6$, therefore explaining the slight decrease in the share of wealth in stocks.

Changing the correlation between the factor innovation and the permanent income shock ($\rho_{n,\varepsilon}$) does not materially affect the average share of wealth in stocks (Graphs B and D for

¹¹Campbell and Viceira (1999) quantify hedging demands by comparing hedging demands from a model with a factor predicting returns relative to the myopic model with a constant share of wealth in stocks. We consider the i.i.d. model as the equivalent of the myopic model in our case since portfolios in the i.i.d. model do not exhibit any time variation in response to the factor realizations. For a specific correlation (ρ) we can compute these hedging demands by comparing the two simulated profiles and compute the percentage differences: $hedg(\rho) = 100 * (\alpha_{\text{factor}}(\rho) - \alpha_{\text{I.I.D.}}(\rho)) / \alpha_{\text{I.I.D.}}(\rho)$. For space considerations we only report the simulated shares of wealth in stocks generated from different correlations of interest.

the two cases), and therefore also does not substantially affect average wealth accumulation.

What happens when the correlation between the permanent earnings shock and the stock market innovation is raised from 0.15 to 0.5? Figure 5 plots the average life cycle portfolio share allocation for the mean reversion model versus the i.i.d. model when this correlation is 0.5. There is a substantial difference across the life-cycle profiles both for the baseline case (Graph A) and the higher factor variability case (Graph B). Interestingly, the average share of wealth in stocks does not drop as much in the mean reversion model as it does in the i.i.d. model. This happens because the higher correlation does not push the share of wealth in stocks to zero for the highest factor realizations that signal high future expected stock returns.

V Model Parameter Uncertainty

Model uncertainty is a feature in the stock market predictability literature as emphasized, for example, in Xia (2001), Brandt et. al. (2005) and Pastor and Stambaugh (2012). We do not introduce learning in the model because this can be a substantial extension that requires either Kalman filtering or Bayesian updating techniques, extensions that we view as an interesting and challenging future research project in the context of life cycle models. We therefore take a more simplistic approach and assume that the investor observes a noisy signal of two key structural parameters, whereas the truth is generated with the baseline parameters analyzed above. We focus on the persistence of the factor (ϕ) and the correlation between the stock return and factor innovations ($\rho_{z,\varepsilon}$).

For ϕ , we assume the investor expects three different values: the baseline value (0.9) with

a 50% probability, a lower persistence (0.7) with 25% probability and a higher persistence (0.93) with a 25% probability. This is necessarily asymmetric to respect factor stationarity. For $\rho_{z,\varepsilon}$ we assume the investor expects two different values: the baseline (-0.6) with a 50% probability and an extreme equal to zero with a 50% probability.

The recursive model is modified by simply assuming that the investor is aware of this uncertainty and takes expectations over these three possibilities, weighting each case with the assumed probability. In the case of $\rho_{z,\varepsilon}$, for example, the current value function is the result of an optimal saving-portfolio choice that expects next period returns to be generated with $\rho_{z,\varepsilon} = -0.6$ and with $\rho_{z,\varepsilon} = 0.0$ with a 50% probability each. The value function is not otherwise affected since investors are assumed not to learn the true value of this parameter. We then simulate stock returns assuming a 50-50 realization of $\rho_{z,\varepsilon}$. It turns out that results are not quantitatively affected if instead we simulate by fixing $\rho_{z,\varepsilon}$ at either -0.6 or 0.0 .

Figure 6 shows the results from introducing model uncertainty in these particular ways. Average portfolios are slightly more balanced over the life cycle (Graphs B (for $\sigma_\varepsilon = 0.007$) and D (for $\sigma_\varepsilon = 0.015$)), especially for the model with the correlation uncertainty, generating also slightly lower wealth accumulation as a result (more clearly seen in Graph C). The effect is larger for the case where factor variability is larger (compare Graphs D and B), but the quantitative average effects are not substantial relative to the no uncertainty cases.

VI Are Lifestyle funds Optimal?

A Deviations from the I.I.D. model

Financial advisors argue that the share of wealth in stocks should decrease as the investor approaches retirement and qualitatively this is what the i.i.d. model predicts as well. Nevertheless, we have seen that a factor model will generate substantial variation in the share of wealth in stocks over the life cycle based on the factor realization. The intuitive argument is that households retiring in 2008 when the stock market had lost a substantial percentage of its value should not have followed blindly the rule followed by life style funds.

In this section we evaluate how important this intuition might be. We start a simulation from the beginning of life but assume that the investor follows the asset allocation implied by the same factor state throughout the lifecycle. This could arise from heterogeneous expectations about the underlying factor among different investors. As we have seen before, this implies that households faced with the lower factor realization should have lower exposure in the stock market. Figure 7 produces such a diagram by assuming different factor realizations starting from each of the 15 different initial states that persist throughout the life cycle. Graph B produces the same information by comparing the highest and lowest asset allocation experiences relative to the Vanguard recommendation and Graphs C and D repeat the same exercise for the higher factor variability case.

We can observe from these figures that the Vanguard life style recommendation is between the upper and lowest factor realization asset allocation profiles and that the deviation is even more pronounced when the factor is more volatile. Thus, the Vanguard recommendation is roughly correct on average but for an investor with market timing ability or expectations

about time-varying expected risk premia, this recommendation will deviate substantially from the optimal portfolio choice. We next evaluate the welfare losses from ignoring stock market mean reversion.

B Welfare Evaluations

To calculate welfare changes we use the value functions across different experiments. Given that we have solved for saving, portfolio choices and value functions for all periods in the life cycle, we know that the value functions at a particular age are a sufficient statistic for welfare effects. Let $v_o(x_{it}, f_t)$ be the value function for the benchmark model and $v_n(x_{it}, f_t)$ be the value function for a new model. We compute a measure of welfare change for a particular age group (*age*) as:

$$(7) \quad \bar{\mu}_{\text{age}} = \text{average of} \left[\left(\frac{v_n(x_{it}, f_t)}{v_o(x_{it}, f_t)} \right)^{\frac{1}{1-\gamma}} - 1 \right] \text{ for all } i \in I_{\text{age}} \text{ and all factor states.}$$

This is the unconditional (across factor states) certainty consumption equivalent because we convert the change of the value into the dimension of expenditure before taking the average.

Figure 8 plots the life cycle certainty equivalents in percent when returns are simulated based on the mean reversion model and the comparison is between the mean reversion and the i.i.d. model and between the mean reversion and Vanguard recommendation. Graph A reports the results for the baseline model ($\sigma_\varepsilon = 0.007$) and Graph B for the higher factor volatility model ($\sigma_\varepsilon = 0.015$). Graph A illustrates substantial welfare losses from following the i.i.d. model relative to the optimal portfolio rule in the presence of the factor, and the welfare losses are even more substantial when following the Vanguard recommendation. This

arises naturally given that the deviations of the average portfolio allocations are even larger between the Vanguard recommendation and the factor model than the ones between the i.i.d. and factor model (Graph B of Figure 2). With higher factor volatility (Graph B) these effects are magnified as the differences across the average portfolio shares further diverge (Graph D of Figure 2).

We make two observations based on the results in Figure 8. First, the welfare losses are economically significant: they can reach 2% to 7% of consumption equivalents depending on the specification, and this represents a substantial welfare loss. Second, the losses tend to get maximized at around age 50, whereas average wealth accumulation is maximized at the exogenous retirement age (65). What can explain these findings?

To better understand these welfare shapes and magnitudes, it is helpful to recall that given the preference for consumption smoothing, welfare in the model is increasing in consumption and decreasing in the variability of consumption. We can therefore gain an insight on where the welfare differences are coming from by investigating how average consumption and consumption inequality evolve over the life cycle across the three models. To do so we compute the average consumption and standard deviation of consumption as a cohort ages. Mean consumption is very similar over the life cycle across the three models in a life cycle graph. To emphasize the differences we therefore report the percentage difference between the profiles generated by the baseline versus the i.i.d. model and the baseline versus the Vanguard recommendation. These differences are plotted in Figure 9 for the baseline case ($\sigma_\varepsilon = 0.007$ for Graphs A and B) and for the higher factor volatility case (Graphs C and D).

The Vanguard recommendation generates substantially more volatility in consumption over the working part of the life cycle (Graphs B and D), despite generating mildly higher

average consumption in the baseline case (Graph A). Given the preferences for smoother consumption, this increased consumption inequality translates into a welfare loss that essentially gets maximized at mid life (around age 50), justifying the peak in welfare loss depicted in Figure 8.

The i.i.d. model on the other hand generates lower mean consumption but lower volatility over the life cycle. The biggest welfare loss is therefore coming from lower mean consumption and this is even more pronounced in the case where the factor is more volatile (compare Graph C with Graph A, for example). Consumption variability is actually lower with the i.i.d. model since portfolio rules are a lot more stable, and this is especially so for the more volatile factor (Graph D relative to Graph B).

The question remains as to why welfare losses are maximized at around age 50. This happens because the saving rate (defined as the percent of labor income that is saved) turns negative at around that age. This is driven by the hump shape in the labor income process and the fact that the rate of average wealth accumulation begins to slow down after that age as average labor income begins to fall. Given the reduction in the saving rate, the welfare loss from following an imperfect portfolio rule is reduced.

C Out-of-Sample Analysis

We next compare how realized wealth and certainty equivalents evolve out of sample, partly adapting the calculations in Lan (2015) for our model. Specifically, for the two cases (high and low σ_ε) we start simulating for every age using the realized dividend yield and stock returns between 1985 and 2014 and between 1999 and 2014. For instance, given a simulated initial distribution of assets per age in 1985, from 1986 onwards we use the realized dividend

yield to pick the relevant factor state and the realized stock return to shock financial asset returns. We save financial wealth for every age group between 1986 and 2014, as well as the CEQs as defined in (7). For an investor at age 30, for example, we track the evolution of individual wealth and CEQs over the 1986-2014 period and average in the cross section every year.

It might be helpful to start the discussion by plotting the wealth levels for different models. Figure 10 shows how the mean wealth evolves for different starting ages according to the different models. As can be seen, the baseline models typically generate higher mean wealth than either the i.i.d. or Vanguard rule would predict. Moreover, younger households display the mean wealth accumulation rising over time (and over the life cycle), while older households can display the decreasing wealth after retirement.

Tables 1 and 2 report the mean differences across CEQs for different models ($\sigma_\varepsilon = 0.007$ and $\sigma_\varepsilon = 0.015$, respectively). Specifically, as in Lan (2015), the mean differences across CEQs for each age are computed over the 1985 to 2014 (Panel A) and 1999 to 2014 (Panel B) periods. The t-statistics are computed using a Newey-West procedure to correct for serial correlation. One conclusion from Table 1 is that the increase in welfare is not always statistically significant (compare the baseline model with the i.i.d. stock returns model in the 1985 to 2014 period). Nevertheless, the baseline model statistically and economically outperforms the Vanguard model in both periods, while the baseline model statistically and economically outperforms both the i.i.d. and Vanguard models in the 1999 to 2014 period. These conclusions are strengthened when the factor is more volatile (Table 2) since in that instance the dividend yield as a signal generates more aggressive investor behavior. These results emphasize the importance of searching for good signals when devising long term asset

allocation strategies, but also how the same variable can generate different conclusions over different time periods.

D Rules of Thumb

Can a rule of thumb be devised in the presence of factor predictability that dominates widely used recommendations like the Vanguard one? Based on our results we devise approximate asset allocation rules that are similar to Vanguard's but are conditional on the factor realization for the baseline case: they are more aggressive for higher factor realizations and less aggressive for lower ones. For the baseline case, and relative to the Vanguard recommendation, the following different rules of thumb are applied. If the factor realization is above 5%, the investor allocates all financial wealth to the stock market. If the factor realization is between 2 and 5%, then the Vanguard rule is applied. If the factor realization is below 2%, then the maximum of zero and the Vanguard rule minus 50% is chosen. We also report results for changing the thresholds between 1% and 6%, rather than between 2% and 5%, and also by adjusting the allocation relative to the Vanguard recommendation by 30% instead of 50%. The consumption/saving rules are chosen optimally in both cases. We think a richer model of retirement uncertainties is needed to make similar recommendations for the retirement period and we leave this extension for future work.

Tables 3 and 4 report the mean differences across CEQs for the two models ($\sigma_\varepsilon = 0.007$ and $\sigma_\varepsilon = 0.015$, respectively). A positive percentage means that the rule of thumb dominates the Vanguard recommendation. We can see that over most time periods and experiments there is a statistically and economically significant welfare improvement relative to the Vanguard recommendations.

We can also provide a rule of thumb of the saving rate (the proportion of labor income that is saved). Depending on when the household begins saving for retirement, the model predicts a rule of thumb of around 10% of labor income saved for thirty years of working life to achieve adequate consumption smoothing, regardless of the factor variability. Implementing these rules of thumb lowers the welfare loss relative to both the i.i.d. and Vanguard models, assuming stock returns are generated through the mean reversion process but we do not report these welfare comparisons due to space considerations.

VII Conclusion

In the presence of stock market predictability, undiversifiable labor income risk and exogenously imposed liquidity constraints, the consumption policy rule has a similar shape with consumption functions derived in the buffer stock saving literature. Optimal portfolio choice is shown to be heavily dependent on the realization of the factor predicting future returns. In the baseline case where the factor is always positive, the share of wealth in stocks is a parallel shift of the i.i.d. model, and the extent of the shift depends on the magnitude of the factor realization and the degree that short sale constraints bind. When the factor is more volatile and can take negative expected values, portfolio holdings will very often be either completely allocated in the stock market or in the riskless asset market. The large welfare losses from failing to condition on market information support the case for enhanced TDFs (ETDFs) that condition asset allocation on expected risk premia.

Future directions of research include the explicit introduction of tax deferred retirement accounts (for the i.i.d. case, see Gomes, Michaelides, and Polkovnicheno (2009)), an explicit

learning mechanism about the true underlying model through either a Kalman filtering or Bayesian learning approach (Brandt et. al. (2005) and Pastor and Stambaugh (2012)), an explicit treatment of housing and introducing time-varying volatility and risk aversion (through a stochastic discount factor, for example). All these extensions will require additional computational power to achieve the desired required solution accuracy but will further improve our understanding of life cycle portfolio choice under uncertainty and offer scientific advice to billions of households increasingly making their own individual financial decisions.

References

- Angerer, X., and P. Lam. “Income Risk and Portfolio Choice: An Empirical Study.” *Journal of Finance*, LXIV (2009), 1037–1055.
- Balduzzi, P., and A. W. Lynch. “Transaction Costs and Predictability: Some Utility Cost Calculations.” *Journal of Financial Economics*, 52 (1999), 47–78.
- Barberis, N. “Investing for the Long Run when Returns are Predictable.” *Journal of Finance*, 55 (2000), 225–264.
- 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.
- Bonaparte, Y.; G. Korniotis; and A. Kumar. “Income Hedging and Portfolio Decisions.” *Journal of Financial Economics*, 113 (2014), 300–324.
- Brandt, M. “Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach.” *Journal of Finance*, 54 (1999), 1609–1646.

- Brandt, M.W.; A. Goyal; P. Santa-Clara; and J.R. Stroud. “A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning about Return Predictability.” *Review of Financial Studies*, 18 (2005), 831–873.
- Brennan, M.; E. Schwartz; and R. Lagnado. “Strategic Asset Allocation.” *Journal of Economic Dynamics and Control*, 21 (1997), 1377–1403.
- Brennan, M., and Y. Xia. “Dynamic Asset Allocation under Inflation.” *Journal of Finance*, 57 (2002), 1201–38.
- Campbell, J. Y., and L. Viceira. “Consumption and Portfolio Decisions When Expected Returns are Time Varying.” *Quarterly Journal of Economics*, 114 (1999), 433–495.
- 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.
- Campbell, J.; Y. L. Chan; and L. Viceira. “A Multivariate Model of Strategic Asset Allocation.” *Journal of Financial Economics*, 67 (2003), 41–80.
- Carroll, C. “Buffer Stock Saving and the Life Cycle / Permanent Income Hypothesis.” *Quarterly Journal of Economics*, CXII (1997), 3–55.
- Cocco, J., and P. Lopes. “Reverse Mortgage Design.” LBS/LSE mimeo, (2015).
- Cocco, J.; F. Gomes; and P. Maenhout. “Consumption and Portfolio Choice over the Life–Cycle.” *The Review of Financial Studies*, 18 (2005), 491–533.

Cooper, R., and G. Zhu. “Household Finance over the Life-Cycle: What does Education Contribute?” Penn State working paper (2014).

Davidoff, T. “Can High Costs Justify Weak Demand for the Home Equity Conversion Mortgage?” *Review of Financial Studies*, 28 (2015), 2364–2398.

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.

Deaton, A. “Saving and Liquidity Constraints.” *Econometrica*, 59 (1991), 1221–1248.

Donaldson, S.; F. Kinniry; R. Aliaga-Diaz; A. Patterson; and M. DiJoseph. “Vanguard’s Approach to Target-Date Funds.” (2013).

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.

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

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

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

Haliassos, M., and A. Michaelides. “Portfolio Choice and Liquidity Constraints.” *International Economic Review*, 44 (2003), 144–177.

- Heaton, J., and D. Lucas. “Portfolio Choice in the Presence of Background Risk.” *The Economic Journal*, 110 (2000), 1–26.
- Kim, T. S., and E. Omberg. “Dynamic Nonmyopic Portfolio Behavior.” *Review of Financial Studies*, 9 (1996), 141–161.
- 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.
- Lan, C. “An Out-of-Sample Evaluation of Dynamic Portfolio Strategies.” *Review of Finance*, 19 (2015), 2359–2399.
- Lynch, A., and S. Tan. “Labor Income Dynamics at Business-Cycle Frequencies: Implications for Portfolio Choice.” *Journal of Financial Economics*, 101(2011), 333–359.
- Munk, C., and C. Sorensen. “Dynamic asset allocation with stochastic income and interest rates.” *Journal of Financial Economics*, 96 (2010), 433–462.
- Pastor, L., and R. F. Stambaugh. “Are Stocks Really Less Volatile in the Long Run?” *The Journal of Finance*, LXVII (2012), 431–477.
- Pettenuzzo, D.; A. Timmermann; and R. Valkanov. “Forecasting Stock Returns Under Economic Constraints.” *Journal of Financial Economics*, 114 (2014), 517–553.
- Tauchen, G. “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters*, 20 (1986), 177–181.
- Tauchen, G., and R. Hussey. “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models.” *Econometrica*, 59 (1991), 371–396.

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

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

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.

Weil, P. “Non-Expected Utility in Macroeconomics.” *Quarterly Journal of Economics*, 2 (1990), 29–42.

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

FIGURE 1

Policy Function Comparison

Figure 1 presents the policy function comparisons between the benchmark stock market mean reversion results and the I.I.D. stock returns case.

FIGURE 2

Life-Cycle Profile Comparison

Figure 2 presents the life-cycle profile comparison between the benchmark stock market mean reversion results, the Vanguard recommendation and the I.I.D. stock returns model.

FIGURE 3

Life-Cycle Individual Share of Wealth in Stocks

Figure 3 presents the life-cycle individual share of wealth in stocks for different initial factors starting from age 21 for the benchmark stock market mean reversion model.

FIGURE 4

Life-Cycle Profiles for Average Wealth and Portfolio Shares for Different Parameters

Figure 4 presents life-cycle profiles for average wealth and portfolio shares for different parameters. The three cases are (i) the benchmark case ($\gamma = 5$ and $\psi = 0.5$) (ii) changing the correlation between permanent labor income shocks and the factor innovations ($\rho_{n,\varepsilon}$) from 0.0 to 0.15 and (iii) changing the correlation between the stock return shocks and factor innovations ($\rho_{z,\varepsilon}$) from -0.6 to 0.0.

FIGURE 5

Life-Cycle Mean Portfolio Shares in Stocks

Figure 5 presents life-cycle mean portfolio shares in stocks. The average life-cycle portfolio shares in stocks for the stock market mean reversion model and the I.I.D. stock returns model when the correlation between the permanent earnings shocks and the stock market innovations is 0.5 instead of 0.15.

FIGURE 6

Life-Cycle Profiles for Average Wealth and Portfolio Shares for Different Parameters

Figure 6 presents the life-cycle profiles for average wealth and portfolio shares for different parameters. The three cases are (i) the benchmark case ($\gamma = 5$ and $\psi = 0.5$) (ii) the benchmark case with uncertain ϕ and (iii) the benchmark case with uncertain $\rho_{z,t}\varepsilon$.

FIGURE 7

Portfolio Shares with Different Constant Factor Realizations

Figure 7 presents average portfolio shares with different constant factor realizations for the stock market mean reversion model. Graphs B and D report the highest and lowest states and also report the Vanguard recommendation that lies in between the two extremes.

FIGURE 8

Welfare Evaluation

Figure 8 presents average consumption certainty equivalents in percent. We assume returns are generated based on the mean reversion model but households are either using the I.I.D. model or the Vanguard policy rule. Welfare loss is relative to the baseline case.

FIGURE 9

Consumption Evaluation

Figure 9 presents percentage changes in average consumption and standard deviation of consumption over the life cycle. The two cases are (i) the I.I.D. case ($\gamma = 5$ and $\psi = 0.5$) relative to the baseline model and (ii) the Vanguard case ($\gamma = 5$ and $\psi = 0.5$) relative to the baseline model.

FIGURE 10

Mean Wealth Profiles for Different Starting Investor Ages in 1985

Figure 10 presents mean wealth profiles for different starting investor ages in 1985. The initial wealth distribution for all ages is generated through a simulation of the baseline model. From that same initial wealth distribution in 1985, investors start behaving according to the baseline, i.i.d. and Vanguard models, while realized dividend yields and stock returns are used from 1986 to 2014 to generate individual wealth profiles.

Table 1

Out of Sample Analysis with $\sigma_\varepsilon = 0.007$

Table 1 presents the mean differences in certainty equivalents (CEQs) between the benchmark model (with $\sigma_\varepsilon = 0.007$) and the I.I.D. stock returns and Vanguard models, respectively. The t-statistics, based on the asymptotic distribution, are in parenthesis. In Panel A, for each starting age and initial wealth distribution based on an initial life cycle simulation, we do the out of sample analysis from 1985 to 2014, using the realized stock returns and the realized dividend yield to pick the relevant policy function. In Panel B, we do the same from 1999 to 2014.

Starting	Panel A: 1985-2014		Panel B: 1999-2014	
Age	I.I.D.	Vanguard	I.I.D.	Vanguard
30	0.3(0.3)	-1.5(-1.6)	-1.5(-5.1)	-4.9(-13.6)
40	-0.3(-0.4)	-0.5(-0.6)	-1.8(-4.7)	-6.3(-12.4)
45	-0.1(-0.1)	0.0(0.0)	-1.8(-4.6)	-6.0(-11.0)
50	0.3(0.2)	0.9(0.8)	-1.6(-4.3)	-5.7(-9.9)
55	0.7(0.5)	-0.1(-0.2)	-1.6(-4.1)	-5.5(-10.8)
60	0.0(-0.0)	-3.3(-11.6)	-1.8(-4.6)	-5.3(-13.9)
65	-0.9(-0.6)	-7.5(-9.5)	-3.0(-5.9)	-4.6(-11.0)
70	-0.6(-0.4)	-12.8(-11.9)	-3.3(-6.5)	-5.2(-64.0)

TABLE 2

Out of Sample Analysis with $\sigma_\varepsilon = 0.015$

Table 2 presents the mean differences in certainty equivalents (CEQs) between the benchmark model (with $\sigma_\varepsilon = 0.015$) and the I.I.D. stock returns and Vanguard models, respectively. The t-statistics, based on the asymptotic distribution, are in parenthesis. In Panel A, for each starting age and initial wealth distribution based on an initial life cycle simulation, we do the out of sample analysis from 1985 to 2014, using the realized stock returns and the realized dividend yield to pick the relevant policy function. In Panel B, we do the same from 1999 to 2014.

Starting	Panel A: 1985-2014		Panel B: 1999-2014	
Age	I.I.D.	Vanguard	I.I.D.	Vanguard
30	-5.8(-13.5)	-6.2(-16.0)	-6.5(-15.0)	-8.6(-19.6)
40	-6.4(-10.4)	-5.5(-15.7)	-6.7(-14.9)	-9.7(-20.4)
45	-6.6(-8.9)	-5.4(-13.9)	-6.6(-16.4)	-9.6(-18.0)
50	-6.8(-8.4)	-5.3(-10.6)	-6.5(-18.2)	-9.4(-16.9)
55	-7.3(-9.8)	-7.3(-15.7)	-6.5(-16.9)	-9.5(-26.5)
60	-8.0(-10.9)	-10.1(-13.5)	-7.1(-17.8)	-9.8(-29.0)
65	-8.2(-8.6)	-13.5(-16.7)	-9.1(-12.7)	-10.0(-15.9)
70	-9.4(-11.1)	-19.7(-12.6)	-0.1(-13.8)	-10.8(-21.0)

TABLE 3

Out of Sample Analysis with $\sigma_\varepsilon = 0.007$

Table 3 presents the mean differences in certainty equivalents (CEQs) between the Vanguard model (with $\sigma_\varepsilon = 0.007$) and the different rule-of-thumb, respectively. The t-statistics, based on the asymptotic distribution, are in parenthesis. In Panel A, for each starting age and initial wealth distribution based on an initial life cycle simulation, we do the out of sample analysis from 1985 to 2014, using the realized stock returns and the realized dividend yield to pick the relevant policy function. In Panel B, we do the same from 1999 to 2014. U5L2R3: If the dividend yield is greater than 5%, then set the share of wealth in stocks to 100%. If the dividend yield is below 2%, then take the maximum of the Vanguard recommendation minus 30% and zero. Otherwise, follow the Vanguard recommendation. U5L2R5: Same as U5L2R3 but subtract 50% from Vanguard recommendation when the dividend yield is below 2%. U6L1R3: Same as U5L2R3 but the thresholds of inaction are between 1% and 6% instead of 2% and 5%. U6L1R5: Same as U6L1R3 but subtract 50% from Vanguard recommendation when the dividend yield is below 2%.

Starting Age	Panel A: 1985-2014					Panel B: 1999-2014				
	U5L2R3	U6L1R3	U5L2R5	U6L1R5		U5L2R3	U6L1R3	U5L2R5	U6L1R5	
30	1.5(4.5)	0.8(21.9)	1.8(4.2)	0.8(20.6)		1.7(19.2)	0.8(95.6)	0.0(15.2)	0.9(109.1)	
40	1.4(3.1)	0.7(61.3)	1.7(3.1)	0.7(69.6)		2.5(12.1)	0.8(48.5)	3.4(9.91)	0.9(32.2)	
45	1.3(3.3)	0.7(9.4)	1.6(3.2)	0.7(9.6)		2.8(10.2)	0.8(32.5)	3.7(8.4)	0.8(49.8)	
50	1.7(3.2)	0.6(6.3)	2.4(2.5)	0.6(6.1)		3.0(8.6)	0.8(44.2)	4.0(7.3)	0.8(71.7)	
55	1.6(3.5)	0.5(11.2)	1.5(3.1)	0.5(11.1)		3.1(8.0)	0.7(172.4)	4.2(7.2)	0.7(38.4)	
60	1.0(3.2)	0.5(3.2)	0.9(2.9)	0.5(3.24)		3.5(10.2)	0.7(20.5)	4.7(8.4)	0.7(18.7)	
65	1.2(3.7)	0.8(4.2)	1.2(3.6)	0.9(4.30)		4.4(10.5)	0.9(39.7)	6.3(8.5)	0.9(38.4)	
70	2.0(6.5)	1.4(7.6)	2.1(6.8)	1.4(8.1)		4.9(10.9)	1.2(66.1)	5.3(10.5)	1.3(57.0)	

TABLE 4

Out of Sample Analysis with $\sigma_\varepsilon = 0.015$

Table 4 presents the mean differences in certainty equivalents (CEQs) between the Vanguard model (with $\sigma_\varepsilon = 0.015$) and the different rule-of-thumb, respectively. The t-statistics, based on the asymptotic distribution, are in parenthesis. In Panel A, for each starting age and initial wealth distribution based on an initial life cycle simulation, we do the out of sample analysis from 1985 to 2014, using the realized stock returns and the realized dividend yield to pick the relevant policy function. In Panel B, we do the same from 1999 to 2014. U5L2R3: If the dividend yield is greater than 5%, then set the share of wealth in stocks to 100%. If the dividend yield is below 2%, then take the maximum of the Vanguard recommendation minus 30% and zero. Otherwise, follow the Vanguard recommendation. U5L2R5: Same as U5L2R3 but subtract 50% from Vanguard recommendation when the dividend yield is below 2%. U6L1R3: Same as U5L2R3 but the thresholds of inaction are between 1% and 6% instead of 2% and 5%. U6L1R5: Same as U6L1R3 but subtract 50% from Vanguard recommendation when the dividend yield is below 2%.

Starting Age	Panel A: 1985-2014					Panel B: 1999-2014				
	U5L2R3	U6L1R3	U5L2R5	U6L1R5	U5L2R3	U6L1R3	U5L2R5	U6L1R5	U5L2R3	U6L1R5
30	6.1(15.2)	5.0(32.4)	6.7(10.9)	5.2(32.1)	6.4(34.9)	5.1(202.8)	7.4(23.7)	5.3(181.5)		
40	5.7(18.2)	4.6(27.6)	6.2(11.1)	4.6(29.7)	7.5(69.7)	5.2(26.8)	9.0(28.8)	5.4(24.7)		
45	5.9(18.3)	4.2(11.6)	6.7(8.4)	4.2(11.6)	7.8(52.7)	5.1(24.4)	9.4(24.6)	5.2(21.8)		
50	7.0(9.3)	3.8(9.1)	8.1(5.6)	3.8(9.0)	8.0(39.4)	5.1(28.0)	9.7(20.7)	5.1(23.9)		
55	7.3(10.0)	3.4(11.4)	7.2(8.7)	3.5(11.5)	8.3(27.7)	5.0(27.7)	10.3(17.6)	5.0(25.9)		
60	8.7(12.8)	3.4(19.2)	8.5(12.2)	3.4(19.8)	9.5(21.5)	5.0(22.5)	12.2(12.6)	4.9(20.3)		
65	12.4(14.4)	4.3(9.8)	12.1(13.8)	4.3(9.9)	11.7(14.1)	5.6(121.6)	14.6(9.8)	5.4(164.0)		
70	18.7(11.6)	6.4(9.2)	18.6(11.4)	6.4(9.4)	13.1(14.4)	6.9(121.1)	13.4(13.6)	6.8(142.4)		