Optimal Portfolio Strategies in the Presence of Regimes in Asset Returns*

Carlos Heitor Campani†, René Garcia‡ and Marcelo Lewin§

This Version Revised on December 2, 2020

Abstract

This paper analyzes optimal portfolio and consumption strategies in a regime-switching economy with unobservable states and predictability of risky asset returns. We develop approximate analytical solutions to the unconstrained dynamic problem. The approximation is shown to be fast and accurate in a four-regime setting with an allocation to four assets compared to the numerical solution developed in Guidolin and Timmermann (2007). The computation time of the approximate solution is shown to be practically independent of the number of assets when no predictors are present and only marginally affected by the number of predictors. While the portfolio policy strongly depends on the current state of the economy, the consumption-to-wealth ratio is roughly state-independent. Predictability considerably changes the optimal portfolios. Hedging demands are negligible with regimes and no predictability, but are important with predictability. On the other hand, the consumption-to-wealth ratio is not very impacted by the predictor. We provide an out-of-sample statistical assessment of the returns provided by a multi-regime strategy with respect to a single-regime and to a 1/N strategy.

JEL classification: G11, C02

Keywords: Dynamic asset allocation, Stochastic differential utility, Consumption and portfolio optimal strategies, Regime switching economy, Predictability, Large and small caps, Size effects.

*This paper includes an online appendix in which we provide detailed derivations and extensions at `http://ssrn.com/abstract=BLABLABLA`.

†COPPEAD Graduate School of Business (Federal University of Rio de Janeiro) Research Associate at Edhec-Risk Institute, carlos.heitor@coppead.ufrj.br

‡Université de Montréal and Toulouse School of Economics, rene.garcia@umontreal.ca (Corresponding Author)

§COPPEAD Graduate School of Business (Federal University of Rio de Janeiro), marcelo.lewin@coppead.ufrj.br

1 Introduction

To allocate their wealth between several asset classes over time, investors have to determine how asset returns evolve dynamically. They could assume simple processes with constant coefficients but they will be at odds with the numerous changes that affect their means, volatilities and correlations. In the global financial crisis of 2008-2009, stock returns exhibited drastic changes in their time-series characteristics. More generally, they evolve differently in the various phases of the business cycle. Since the seminal paper of Hamilton (1989), Markov-switching models have been widely used to capture these sudden changes in financial markets (see a survey by Ang and Timmermann (2012)). Investors face a challenge since these regimes are not observable with certainty and they need to infer their probabilities of occurrence and their duration based on observed returns and possibly some predictive variables. Moreover, they need to use time-consuming numerical techniques to build optimal dynamic portfolios. In this paper we provide approximate analytical solutions to these complex portfolio allocation problems. Our method is fast and accurate compared to numerical solutions. Adding assets does not change computation time while adding predictors only affects it marginally. Therefore, our method offers practical solutions to portfolio managers that have to rebalance sizable portfolios frequently.

We assume that investors solve a general finite-horizon portfolio allocation problem with consumption, in which we distinguish risk aversion and intertemporal substitutability with a stochastic differential utility. We allow for the presence of a predictor variable p_t in a multi-regime economy. The regimes are hidden, and investors have to estimate the probability π_t of each state at each point in time. Our general solution admits a wealth-separable solution of the form:

$$V\left(W_t, p_t, \pi_t, \tau\right) = H\left(p_t, \pi_t, \tau\right) \frac{W_t^{1-\gamma}}{1-\gamma}, \tag{1}$$

in which W_t represents the investor's wealth at time t and $\tau = T - t$ is the time remaining till the final horizon T. We find an approximate linear expression for $H\left(p_t, \pi_t, \tau\right)$ in terms of the state variables p_t, π_t and their squares. The coefficients multiplying the state variables are horizon-dependent and solve a system of ordinary differential equations.

The optimal portfolio strategy contains the usual myopic allocation and two hedging demands related to the predictor and to the regime probabilities, all inversely proportional to the coefficient of risk aversion. The three components depend also implicitly on the elasticity of intertemporal substitution (EIS). Optimal consumption is affected by the state variables through the function $H(p_t, \pi_t, \tau)$ and is directly impacted by the EIS.

To assess the accuracy of our solution method, we use a model similar to Guidolin and Timmermann (2007), with four regimes in the returns of large and small stocks and bonds. We compare our approximate solution to their optimal numerical solution based on Monte Carlo methods. We show that the approximate portfolio shares are very close to the optimal ones when regimes are assumed to be known or when the investor needs to estimate the probabilities of being in each of the states. We also measure the wealth equivalent utility loss created by using the approximate solution instead of the optimal solution obtained by simulation and conclude that it is negligible.

To evaluate the economic importance of considering multiple regimes we conduct an out-of-sample exercise to compare the returns of a portfolio allocation to large stocks, small stocks, bonds and a risk-free asset for a four-regime model, a single-regime model and a $1/N$ strategy. We assume an investment horizon of 10 years with a monthly rebalancing of the portfolio. We compute the differences in annualized certainty equivalent returns between the four-regime portfolio and the two benchmark portfolios and their bootstrap confidence intervals.

When stock and bond returns are predictable by the dividend yield, we observe that the share invested in large stocks increases in the short run with respect to the case without predictability but mean reversion in the predictor makes it go down after 3 to 4 years up to the 10-year horizon. Small stocks pick up some of the slack left by large stocks but the overall share in stocks still exhibits a mean-reverting pattern. The presence of hedging demands with respect to the predictor increases the expected loss suffered by a myopic investor at long horizons. This contrasts with negligible hedging demands in the absence of a predictor. Accounting for intermediate consumption in the investor's problem, we conclude that the consumption-to-wealth ratio does not vary much with the predictor and with the regimes.

Several papers have studied dynamic optimal portfolio problems under regime-switching processes in discrete time when regimes are unobservable. Ang and Bekaert (2002) solved numerically a two-regime portfolio problem with international assets, while Guidolin and Timmermann (2007) use Monte Carlo techniques to solve an allocation between large and small stocks and bonds in a four-state regime switching model. In continuous time, Liu (2011) examines a consumption-portfolio problem in which the expected return of a single risky asset follows a hidden Markov chain, under ambiguity. The investor's optimal choices are not in closed-form, but characterized in terms of Malliavin derivatives and stochastic integrals. This model nests the expected utility model of Honda (2003).

Many papers analyze the problem of dynamic portfolio allocation with fully observable regimes. Yin and Zhou (2003) find an explicit solution when an investor minimizes the variance of a given fixed expected terminal wealth and does not consume. Sotomayor and Cadenillas (2009) study the consumption-portfolio problem of a power utility investor who maximizes the expected total discounted utility of consumption and find optimal exact portfolio and consumption policies. All market parameters (interest rate, stock drift and volatility) are constant within a regime. They work under infinite horizon and hence are not able to capture and analyze horizon effects.

Lim and Watevai (2012) consider the problem of optimal asset allocation for a regime-switching model with jumps. They find semi-explicit expressions for the optimal policy in terms of the solution of a system of ordinary differential equations with CRRA utility. Çanakoğlu and Özekici (2012) also find explicit solutions for optimal portfolio policy for HARA utilities in terms of solutions of ordinary differential equations. López and Serrano (2015) find closed-form expressions of the optimal value function in a pure-jump model for agents with log-utility and fractional power utility in the case of two-state Markov chains. Xing (2017) and Kraft et al. (2017) solve optimal portfolio and consumption problems with Epstein-Zin recursive utility without regime switching.

Several other papers are related to our work. Graflund and Nilsson (2003) study the relevance of intertemporal hedging and regimes under a dynamic portfolio problem with no consumption in which the investor maximizes power utility from terminal wealth. The discrete-time setting includes one risky asset and a riskless bond, in which the investor rebalances the portfolio monthly. The problem is again solved numerically, conditional on the current observable regime, with the number of regimes being determined with a Monte Carlo likelihood ratio test. They conclude that ignoring the regimes for long-horizon investors is costly, while intertemporal hedging is present in some regimes but not in others.

The importance of regime shifts for modeling asset returns has been examined by a number of researchers. For example, Ang and Bekaert (2004) show the importance of regime switching models in tactical asset allocation. Similarly, Tu (2010) concludes that certainty-equivalent losses associated with ignoring regime switching in portfolio decisions are generally above 2% per year. Guidolin and Hyde (2012) show that vector autoregressive models cannot capture regime shifts in asset returns. They use a three-regime model in which the investor has power utility. All these papers are set in discrete-time and solve the problem numerically, via Monte Carlo methods.

Guidolin and Timmermann (2006, 2008a) use a discrete-time regime switching model for asset returns to solve the portfolio problem (with no consumption), in which the investor has utility over moments of the terminal wealth distribution. The current regime is observable and the approximate optimal portfolio weights are found as the roots of a system of polynomial equations (first-order conditions). Under unobservable regimes, Guidolin and Timmermann (2005, 2007) specify a four-state regime switching model in discrete-time with a richer risky asset menu (one large-stock index, one small-stock index and a ten-year bond). Their optimal choices, under power time-additive utility, are found numerically, with Monte Carlo techniques. Using the same techniques, Guidolin and Timmermann (2008b) consider an allocation over size and value portfolios.

The rest of the paper is structured as follows. Sections 2 and 3 describe the economy and the investor's problems respectively. The solution and its approximation are explained in Section 4. In Section 5, we compute optimal portfolios in a four-regime model with four assets by numerical methods and assess the accuracy of our approximate solution. We also conduct an out-of-sample exercise and introduce a predictor to show how it affects optimal strategies. The case with intermediate consumption is discussed. Section 6 analyzes the impact of adding assets or predictors on the computation time. We conclude in Section 7. An appendix provides all the necessary details about the approximate Bellman equations and the Monte Carlo simulation procedures to assess the accuracy of the approximate solution.

2 The Economy

We consider an economy with a short-term riskless asset and n different risky assets where the uncertainty affecting the evolution of their returns is governed by a continuous-time Markov chain Y_t and $n+1$ independent Brownian processes (one for each risky asset, $Z_{1,t}$, $Z_{2,t}$, ..., $Z_{n,t}$ and one for a single predictor $Z_{p,t}$). The short-term riskless asset has price M_t at time t and follows the deterministic process $\frac{dM_t}{M_t} = r dt$.[1] We first describe the Markov process Y_t and then the dynamics of the risky assets and the predictor variable.

[1] The assumption of a constant instantaneous interest rate r is often made for simplification in portfolio models with regime shifts; see, for example, Guidolin and Timmermann (2007), Liu (2011) or Honda (2003). However, Ang and Bekaert (2002) and Guidolin and Timmermann (2008b) consider the case of regime-switching short-term rates that predict future investment opportunities.

2.1 The regime-switching state variable

The regime-switching state variable Y_t is an independent continuous-time Markov chain, right-continuous and admitting only values in $R = \{1, 2, ..., m\}$, where R represents the finite set of m possible regimes in the economy. The regime-switching process Y_t, starting at any given time t_0 on a given state, remains in this state for an exponentially distributed length of time, and then jumps to another state.

More precisely, considering the current state is i, the probability of jumping to another state j over the next Δt time period is $P_{ij,\Delta t} = \frac{\lambda_{ij}}{\sum_{k\neq i} \lambda_{ik}} \left(1 - e^{-\sum_{k\neq i} \lambda_{ik}\Delta t}\right)$, with $j \neq i \in R$ and $\lambda_{ij} \geq 0$. [2] We define $\lambda_{ii} = -\sum_{k\neq i} \lambda_{ik} \leq 0$ such that $P_{ij,\Delta t} = \frac{\lambda_{ij}}{-\lambda_{ii}} \left(1 - e^{\lambda_{ii}\Delta t}\right)$. The probability of staying in the same regime i over the next Δt time period is therefore given by $P_{ii,\Delta t} = e^{\lambda_{ii}\Delta t}$.

The parameters λ_{ij} ($j \neq i \in R$) are assumed to be constant and represent the density of transition probabilities from regime i to regime j – (as $Lim_{\Delta t\to 0} \frac{P_{ij,\Delta t}}{\Delta t} = \lambda_{ij}$). The closer λ_{ii} is to zero, the more persistent regime i is. When $\lambda_{ii} = 0$, once the economy jumps to state i, it will remain there forever. We also adopt the standard assumption that the inter-regime times are independent, and are also independent of Brownian motions governing the risky assets and the predictor.

2.2 The risky assets and the predictor

The dynamics of the n risky assets is given by:

$$\underbrace{\begin{bmatrix} \frac{dS_{1,t}}{S_{1,t}} \\ \frac{dS_{2,t}}{S_{2,t}} \\ ... \\ \frac{dS_{n,t}}{S_{n,t}} \end{bmatrix}}_{\frac{\mathbf{dS_t}}{\mathbf{S_t}}} = \underbrace{\begin{bmatrix} \mu_{1,t} \\ \mu_{2,t} \\ ... \\ \mu_{n,t} \end{bmatrix}}_{\mu_{s,t}} dt + \underbrace{\begin{bmatrix} \sigma_{11,t} & 0 & ... & 0 \\ \sigma_{21,t} & \sigma_{22,t} & ... & 0 \\ ... & ... & ... & 0 \\ \sigma_{n1,t} & \sigma_{n2,t} & ... & \sigma_{nn,t} \end{bmatrix}}_{\sigma_{s,t}} \underbrace{\begin{bmatrix} dZ_{1,t} \\ dZ_{2,t} \\ ... \\ dZ_{n,t} \end{bmatrix}}_{\mathbf{dZ_t}}, \tag{2}$$

where $Z_{1,t}$, $Z_{2,t}$, ..., $Z_{n,t}$ are standard and independent Brownian motion processes, $\mu_{s,t}$ is a $n \times 1$ column vector of instantaneous expected risk-premia and $\sigma_{s,t}$ is an $n \times n$ lower triangular time-varying volatility matrix. The characterization of the time variation will be made precise in equations (9), (11) and (12).

We also assume that there exists a predictor variable p_t of the risk premia of the risky assets

[2] To be concise, $P_{ij,\Delta t} = P(Y_{\Delta t} = j | Y_t = i), \forall t \in [0, \Delta t]$.

that follows a mean-reverting Ornstein-Uhlenbeck diffusion process:

$$dp_t = \mathbf{D_{p,t}} dt + \sigma_\mathbf{p} \underbrace{\begin{bmatrix} dZ_{1,t} & dZ_{2,t} & \dots & dZ_{n,t} & dZ_{p,t} \end{bmatrix}^T}_{\mathbf{dZ_t^*}}, \tag{3}$$

where the vector $\mathbf{D_{p,t}}$ collects the different predictor means conditional on the regimes:

$$\mathbf{D_{p,t}} = \kappa_p \begin{bmatrix} \bar{p}_1 - p_t & \bar{p}_2 - p_t & \dots & \bar{p}_m - p_t \end{bmatrix}. \tag{4}$$

The parameter κ_p denotes the speed of mean-reversion (which is not regime-dependent) and $\sigma_\mathbf{p} = \begin{bmatrix} \sigma_{p1} & \sigma_{p2} & \dots & \sigma_{pn} & \sigma_{pp} \end{bmatrix}$ is a constant row vector with $n+1$ elements, while $Z_{p,t}$ is another standard Brownian motion, independent from the previously defined Brownian motions. We concatenate the risky assets and predictor volatility matrices in a $(n+1) \times (n+1)$ lower triangular volatility matrix σ_t: $\sigma_t = \begin{bmatrix} \sigma_{s,t} \\ \sigma_p \end{bmatrix}$ and $\Sigma_t = \sigma_t \sigma_t^T$.

The risky asset drifts are affine functions of the predictor with regime-dependent constants. We collect these drifts in a drift matrix $\mathbf{D}_{s,t}$ of dimension nxm, which is time-varying and dependent on the predictor's value p_t:

$$\mathbf{D_{s,t}} = \underbrace{\begin{bmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,m} \\ a_{2,1} & a_{2,2} & \dots & a_{2,m} \\ \dots & \dots & \dots & \dots \\ a_{n,1} & a_{n,2} & \dots & a_{n,m} \end{bmatrix}}_{\mathbb{A}} + \underbrace{\begin{bmatrix} b_1 \\ b_2 \\ \dots \\ b_n \end{bmatrix}}_{\mathbb{B}} \underbrace{\begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}}_{m \text{ columns}} p_t, \tag{5}$$

for simplicity, we assumed that the parameters of matrix $\mathbb{B}$ are not regime-dependent. There seems to be some support for it in Guidolin and Timmermann (2008b).

The drifts of the risky assets and of the predictor are regime-dependent and therefore a function of the unobservable state variable Y_t.

3 The Investor's Problem

Investors do not observe the drifts of the risky assets and the predictor nor their underlying motions but they can learn about them by observing the asset returns and the predictor values. Investors have therefore a filtering problem to solve to infer the probabilities of being in each regime as well as the probabilities of shifting between regimes. Given these probabilities, they can solve an optimal consumption and portfolio allocation problem.

3.1 The filtering problem

We rely on the model with incomplete information about growth regimes introduced in David (1997), Veronesi (2000), David and Veronesi (2013), and extended recently by Berrada, Detemple and Rindisbacher (2018). Let us define the drift column vector $\mu_{\mathbf{t}} = [(\mu_{\mathbf{s,t}})^{\mathsf{T}}, \mathbf{D_{p,t}}]^{\mathsf{T}}$ that collects the drifts of the n risky assets and the drift of the predictor. Given the Markov economy defined in the last section, this drift vector will follow a continuous-time Markov chain with m regimes.

Based on the available information set $\mathfrak{F}_t$, investors are able to infer regime i probability at time t as: $\pi_{it} = \text{prob}(\text{regime} = i|\mathfrak{F}_t)$. From Wonham (1964), the probabilities $\pi_t = (\pi_{1t}, \dots, \pi_{mt})^{\mathsf{T}}$ follow the process described by David and Veronesi (2013):

$$d\pi_t = \Lambda' \pi_t dt + \Sigma_\pi(\pi_t) d\tilde{Z}_t, \tag{6}$$

where $\Lambda = \{\lambda_{i,j}, i = 1, \dots, m, j = 1, \dots, m\}$ and $\Sigma_\pi(\pi_t)$ the $(m \times (n+1))$ matrix with i^{th} row given by:

$$\sigma_i(\pi_t) = \pi_{it}[\mu_i - \bar{\mu}(\pi_t)]^{\mathsf{T}} \Sigma_t^{-1}, \tag{7}$$

with $\bar{\mu}(\pi_t) = \sum_{i=1}^{n} \pi_{it}\mu_i$. The $(n+1) \times 1$ vector $d\tilde{Z}_t$ is defined as:

$$d\tilde{Z}_t = \Sigma_t^{-1}(\mu_t - \bar{\mu}(\pi_t))dt + dZ_t^*. \tag{8}$$

Suppose now that we want to make the volatilities of the risky assets dependent on the Markov chain: $\sigma_{s,t} = \sigma_{s,i}$, when $Y_i = i$. In this case, in continuous time, the dependence of the volatility on the same Markov chain as the drift makes the Markov chain observable, as shown in proposition 2.1 of Krishnamurthy, Leoff and Sass (2018).

To obtain a model which is both consistent and has good econometric properties, Haussmann and Sass (2004) introduce stochastic volatility in a continuous-time HMM (Hidden Markov Model). They propose a regime-switching model where the volatility depends on an observable diffusion. They show that despite the non-constant volatility the chain is still hidden and they derive its finite-dimensional filtering equation. Krishnamurthy, Leoff and Sass (2018) construct, for a given filtration, the volatility process adapted to this filtration that best approximates the MSM-returns (Markov Switching Models) in the mean squared sense. They propose a so-called filter-based volatility model where: $\sigma_{s,t} = \sum_{i=1}^{m} \sigma_{s,i}\mathbb{P}[Y_i = i|\mathfrak{F}_t^{FB}]$, where $\mathfrak{F}_t^{FB}$ is the filtration generated by the observed returns. Thus, the volatilities are adapted to this observation filtration.

We can therefore rewrite the filtered processes for the risky assets and the predictor conditional to the observed returns and predictor filtration as:

$$\frac{\mathbf{dS_t}}{\mathbf{S_t}} = \mathbf{D_{s,t}}\pi_\mathbf{t} dt + (\mathbf{V}\pi_\mathbf{t})\, \mathbf{d\tilde{Z}_t}, \text{ and} \tag{9}$$

$$dp_t = \mathbf{D_{p,t}}\pi_t dt + \sigma_\mathbf{p} \mathbf{d\tilde{Z}_t^*}. \tag{10}$$

The matrix $\mathbf{V}$ is a 1 x m row vector of matrices:

$$\mathbf{V} = \begin{bmatrix} \sigma_\mathbf{s,1} & \sigma_\mathbf{s,2} & \dots & \sigma_\mathbf{s,m} \end{bmatrix}, \tag{11}$$

where each matrix $\sigma_\mathbf{s,i}, i = 1, \dots, m$ is defined as follows:

$$\sigma_\mathbf{s,i} = \begin{bmatrix} \sigma_{11,i} & 0 & \dots & 0 \\ \sigma_{21,i} & \sigma_{22,i} & \dots & 0 \\ \dots & \dots & \dots & 0 \\ \sigma_{n1,i} & \sigma_{n2,i} & \dots & \sigma_{nn,i} \end{bmatrix}. \tag{12}$$

With this notation, we define:

$$\mathbf{V}\pi = \sum_{i=1}^{m} \sigma_{s,i}\pi_i, \tag{13}$$

3.2 The wealth process and preferences

In this setting the investor consumes and invests in assets. Wealth is the value of the portfolio of assets held by the investor. Let W_t denote wealth at time t. If α_t is the $1 \times n$ vector of the proportions of wealth invested in the risky assets (hence $1 - \alpha_\mathbf{t}\mathbf{1}$ is the proportion invested in the riskless asset) and C_t the amount withdrawn from the portfolio for consumption then wealth evolves according to:

$$
\begin{aligned}
dW_t &= (1-\alpha_t \mathbf{1})\, W_t \frac{dM_t}{M_t} + W_t \alpha_t \frac{\mathbf{dS_t}}{\mathbf{S_t}} - C_t dt \\
&= (1-\alpha_t \mathbf{1})\, W_t r dt + W_t \alpha_t \left[\mathbf{D}_{s,t} \pi_t dt + (\mathbf{V}\pi_t)\, \mathbf{d}\tilde{\mathbf{Z}}_t\right] - C_t dt \\
&= W_t r dt + W_t \alpha_t \left[(\mathbf{D}_{s,t}\pi_t - r\mathbf{1})\, dt + (\mathbf{V}\pi_t)\, \mathbf{d}\tilde{\mathbf{Z}}_t\right] - C_t dt,
\end{aligned}
\tag{14}
$$

where $\mathbf{1}$ represents a column vector of n ones and $\alpha_t = [\alpha_{1,t}\ \alpha_{2,t}\ \ldots\ \alpha_{n,t}]$.

Preferences: Given a time-T finite horizon, the investor's preferences over consumption and terminal wealth (bequest) are represented by a continuous-time recursive utility function, also known as stochastic differential utility function, introduced by Duffie and Epstein (1992):

$$
J_t = E_t \left[\int_{u=t}^{T} f(C_u, J_u)\, du + \frac{W_T^{1-\gamma}}{1-\gamma} \right]. \tag{15}
$$

In equation (15), $f(C, J)$ is a normalized aggregator of current consumption and the continuation utility that takes the following form:

$$
f(C, J) = \frac{\beta}{1-\frac{1}{\psi}} (1-\gamma)\, J \left\{ \left[\frac{C}{[(1-\gamma)\, J]^{\frac{1}{1-\gamma}}} \right]^{1-\frac{1}{\psi}} - 1 \right\}, \tag{16a}
$$

where $\beta > 0$ is a rate of time preference, γ a parameter that controls the investor's risk attitudes over the states of nature, while ψ is related to the investor's consumption choices over time. We will refer to γ as the (relative) risk aversion and ψ as the elasticity of intertemporal substitution.

There are two special cases of the normalized aggregator given by equation (16a): $\psi = \frac{1}{\gamma}$ and $\psi = 1$. The first case is the standard, time-additive power utility function, where log-utility optimal choices are obtained as the limit of the respective optimal choices when $\gamma \to 1$. Note that the resulting formulation may not take the standard power utility form, but will imply the same underlying preferences, and hence the same consumption and asset allocation choices. In the second case, interpreted as the limit when $\psi \to 1$, the aggregator takes the following limit form:

$$
\psi = 1 \qquad \to \qquad f(C, J) = \beta\, (1-\gamma)\, J \ln \left\{ \frac{C}{[(1-\gamma)\, J]^{\frac{1}{1-\gamma}}} \right\}. \tag{16b}
$$

When $\psi = 1$ the investor consumes myopically and ignores investment opportunities. When $\psi < 1$, income effects prevail and better opportunities increase consumption. When $\psi > 1$, substitution effects dominate as the investor is willing to postpone consumption and take advantage of improved investment opportunities.

4 Solving the Problem

4.1 Consumption and portfolio policies

We want to choose the consumption C_t and portfolio weights α_t that maximize the value function $V(W_t, p_t, \pi_t, t) = \sup_{\{C_t,\alpha_t\}} [J(W_t, p_t, \pi_t, t)]$. Given (10) and (14), under Ito's rule, we can write the Bellman equation that represents the recursive problem:

$$0 = \sup_{\{C_t,\alpha_t\}} \left[\begin{array}{c} f(C_t, J_t) + \frac{\partial J_t}{\partial t} + J_w [W_t r + W_t \alpha_t (\mathbf{D_{s,t}} \pi_t - r\mathbf{1}) - C_t] \\ + \frac{1}{2} J_{ww} W_t^2 \alpha_t (\mathbf{V}\pi_t)(\mathbf{V}\pi_t)^{\mathbf{T}} \alpha_t^{\mathbf{T}} + J_p \mathbf{D_{p,t}} \pi_t + \frac{1}{2} J_{pp} \sigma_{\mathbf{p}} \sigma_{\mathbf{p}}^{\mathbf{T}} \\ + \sum_{i=1}^{m} J_{\pi_i} \sum_{j=1}^{m} \lambda_{ji} \pi_{j,t} + \frac{1}{2} \sum_{i=1}^{m} J_{\pi_i \pi_i} \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{i},\pi}^{\mathbf{T}} + J_{wp} W_t \alpha_t (\mathbf{V}\pi_t) \overline{\sigma}_{\mathbf{p}}^{\mathbf{T}} \\ + \sum_{i=1}^{m} J_{w\pi_i} W_t \alpha_t (\mathbf{V}\pi_t) \overline{\sigma}_{\mathbf{i},\pi}^{\mathbf{T}} + \sum_{i=1}^{m} J_{p\pi_i} \sigma_{\mathbf{p}} \sigma_{\mathbf{i},\pi}^{\mathbf{T}} + \sum_{i<j} J_{\pi_i \pi_j} \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{j},\pi}^{\mathbf{T}} \end{array} \right], \tag{17}$$

where $\overline{\sigma}_{\mathbf{p}}^{\mathbf{T}}$ and $\overline{\sigma}_{\mathbf{i},\pi}^{\mathbf{T}}$ represent the transposed vectors of respectively $\sigma_{\mathbf{p}}$ and $\sigma_{\mathbf{i},\pi}$ ($i \in R$) without their last element (the last element will not play a role in the solution since the corresponding Brownian motion has no correlation with the wealth dynamics). The subscripts denote partial derivatives (except t, that refers to the value at time t). Thus, the first-order condition for consumption in the recursive problem is given by:

$$V_w = \frac{\partial f(C_t, V_t)}{\partial C_t} = \beta (1-\gamma)^{\frac{\frac{1}{\psi}-\gamma}{1-\gamma}} V_t^{\frac{\frac{1}{\psi}-\gamma}{1-\gamma}} C_t^{-\frac{1}{\psi}}, \tag{18a}$$

from which we solve for the optimal consumption strategy:

$$C_t = \beta^{\psi} V_w^{-\psi} [(1-\gamma) V_t]^{\frac{1-\psi\gamma}{1-\gamma}}, \tag{18b}$$

which simplifies to $C_t = \beta V_w^{-1} V_t (1-\gamma)$ when $\psi = 1$. Inserting equation (18b) into (16a) and (16b) we can write the recursive aggregator in terms of the sole value function:

$$f(V_t) = \frac{1}{1-\frac{1}{\psi}} \left\{ \beta^{\psi} V_w^{1-\psi} [(1-\gamma) V_t]^{\frac{1-\psi\gamma}{1-\gamma}} - \beta (1-\gamma) V_t \right\} \qquad \psi \neq 1, \tag{19a}$$

$$f(V_t) = \beta (1-\gamma) V_t \ln \left\{ \beta V_w^{-1} [V_t (1-\gamma)]^{\frac{-\gamma}{1-\gamma}} \right\} \qquad \psi = 1. \tag{19b}$$

The first-order condition for the portfolio weights gives the following solution:

$$\begin{aligned}\alpha_{\mathbf{t}} &= \frac{V_w}{-V_{ww}W_t}\left(\mathbf{D_{s,t}}\pi_{\mathbf{t}} - r\mathbf{1}\right)^{\mathbf{T}}\left[\left(\mathbf{V}\pi_{\mathbf{t}}\right)\left(\mathbf{V}\pi_{\mathbf{t}}\right)^{\mathbf{T}}\right]^{-\mathbf{1}} \\ &+ \frac{V_{wp}}{-V_{ww}W_t}\overline{\sigma}_{\mathbf{p}}\left(\mathbf{V}\pi_{\mathbf{t}}\right)^{-\mathbf{1}} + \sum_{i=1}^{m}\frac{V_{w\pi_i}}{-V_{ww}W_t}\overline{\sigma}_{\mathbf{i},\pi}\left(\mathbf{V}\pi_{\mathbf{t}}\right)^{-\mathbf{1}}.\end{aligned} \tag{20}$$

The first term identifies the myopic demand while the second and third components capture the hedging demands that will be held by the investor to hedge against future undesirable movements respectively in the predictor and in the state probabilities.

To obtain the final Bellman equation for the problem under recursive utility, we substitute the optimal consumption and portfolio policies just derived into equation (17):

$$\begin{aligned}0 &= f\left(V_t\right) + \frac{\partial V_t}{\partial t} + V_w W_t r - \beta^{\psi} V_w^{1-\psi}\left[\left(1-\gamma\right)V_t\right]^{\frac{1-\gamma\psi}{1-\gamma}} + V_p \mathbf{D_{p,t}}\pi_t + \frac{1}{2}V_{pp}\sigma_{\mathbf{p}}\sigma_{\mathbf{p}}^{\mathbf{T}} \\ &+ \sum_{i=1}^{m} V_{p\pi_i}\sigma_{\mathbf{p}}\sigma_{\mathbf{i},\pi}^{\mathbf{T}} + \sum_{i,j=1}^{m} V_{\pi_i}\lambda_{ji}\pi_{j,t} + \frac{1}{2}\sum_{i,j=1}^{m}\left(V_{\pi_i\pi_j}\sigma_{\mathbf{i},\pi}\sigma_{\mathbf{j},\pi}^{\mathbf{T}} - \frac{V_{w\pi_i}V_{w\pi_j}}{V_{ww}}\overline{\sigma}_{\mathbf{i},\pi}\overline{\sigma}_{\mathbf{j},\pi}^{\mathbf{T}}\right) \\ &- \frac{1}{2}\frac{V_w^2}{V_{ww}}\left(\mathbf{D_{s,t}}\pi_{\mathbf{t}} - r\mathbf{1}\right)^{\mathbf{T}}\left[\left(\mathbf{V}\pi_{\mathbf{t}}\right)\left(\mathbf{V}\pi_{\mathbf{t}}\right)^{\mathbf{T}}\right]^{-\mathbf{1}}\left(\mathbf{D_{s,t}}\pi_{\mathbf{t}} - r\mathbf{1}\right) - \frac{1}{2}\frac{V_{wp}^2}{V_{ww}}\overline{\sigma}_{\mathbf{p}}\overline{\sigma}_{\mathbf{p}}^{\mathbf{T}} \\ &- \frac{V_w V_{wp}}{V_{ww}}\overline{\sigma}_{\mathbf{p}}\left(\mathbf{V}\pi_{\mathbf{t}}\right)^{-\mathbf{1}}\left(\mathbf{D_{s,t}}\pi_{\mathbf{t}} - r\mathbf{1}\right) - \sum_{i=1}^{m}\frac{V_w V_{w\pi_i}}{V_{ww}}\overline{\sigma}_{\mathbf{i},\pi}\left(\mathbf{V}\pi_{\mathbf{t}}\right)^{-\mathbf{1}}\left(\mathbf{D_{s,t}}\pi_{\mathbf{t}} - r\mathbf{1}\right) \\ &- \sum_{i=1}^{m}\frac{V_{wp}V_{w\pi_i}}{V_{ww}}\overline{\sigma}_{\mathbf{p}}\overline{\sigma}_{\mathbf{i},\pi}^{\mathbf{T}}\end{aligned} \quad , \tag{21}$$

while the final Bellman equation for the case without consumption is given in Appendix A.

The problem admits a wealth-separable solution presented in equation (1), but the Bellman equation cannot be solved analytically due to the presence of nonlinear terms. In the next section, we propose an approximate analytical solution based on a log-linear approximation of the function $H\left(p_t, \pi_{\mathbf{t}}, \tau\right)$ defined in equation (1).

4.2 The approximate analytical solution

Our approximation is inspired by solutions to a similar consumption-portfolio problem with no regimes. In an infinite horizon setting, Campbell and Viceira (1999) and Campbell et al. (2004) provided an approximation for the discrete time case in which the log consumption ratio is approximated around its unconditional long-run mean. Campbell et al. (2004) extended this approximate solution method to a continuous-time setting. The infinite horizon is crucial since the coefficients of the log linear consumption-to-wealth ratio are constant. Campani and

Garcia (2019) develop a solution when the horizon is finite. In this case, the long-run mean for consumption may not be a good proxy for short-horizon investors and thus a good point at which to perform the approximation. Moreover, under finite horizon, the consumption-to-wealth ratio is time-varying even if the state variables do not vary.

They propose to approximate the consumption-to-wealth ratio around a fixed point of the state-variable space. Given the persistence of usual predictors of stock returns, they choose the long-run mean of the predictor. The suggested approximation is:

$$\exp\left(\ln\frac{C_t}{W_t}\right)\approx\exp\left(\ln\frac{C_t}{W_t}\right)_{P_t=\overline{P}}+\exp\left(\ln\frac{C_t}{W_t}\right)_{P_t=\overline{P}}\left[\ln\frac{C_t}{W_t}-\left(\ln\frac{C_t}{W_t}\right)_{P_t=\overline{P}}\right]. \quad (22)$$

They guess a solution of the form: $\ln\frac{C_t}{W_t}=A_1(\tau)+A_2(\tau)P_t+\frac{A_3(\tau)}{2}P_t^2$ and obtain the following expression for the value function:

$$V(W_t,P_t,\tau)=\frac{W_t^{1-\gamma}}{1-\gamma}\exp\left\{(1-\gamma)\left[A_1(\tau)+A_2(\tau)P_t+\frac{A_3(\tau)}{2}P_t^2\right]\right\},$$

with terminal condition $V(W_t,P_t,0)=\frac{W_T^{1-\gamma}}{1-\gamma}$.

By substituting this expression into the Bellman equation, they can solve for $A_i(\tau)$, $i=1,2,3$, in a system of ordinary differential equations with boundary conditions equal to zero since the value function should converge to the power utility of wealth irrespectively of state-variable values at the final horizon.

We follow the same approach to construct our approximate solution. Our state space includes now both the predictor and the regime probabilities. Therefore, since:

$$\frac{C_t}{W_t}=\beta^{\psi}H(p_t,\pi_t,\tau)^{\frac{1-\psi}{1-\gamma}}, \quad (23)$$

we propose the following functional form for the H function:

$$H(p_t,\pi_t,\tau)=\exp\left[\begin{array}{c}A_0(\tau)+A_p(\tau)p_t+\sum_{i=1}^{m}A_i(\tau)\pi_{i,t}+B_p(\tau)p_t^2\\+\sum_{i=1}^{m}B_i(\tau)\pi_{i,t}^2+\sum_{i=1}^{m}C_{pi}(\tau)p_t\pi_{i,t}+\sum_{i<j}C_{ij}(\tau)\pi_{i,t}\pi_{j,t}\end{array}\right]. \quad (24)$$

All coefficients are time-varying and depend on the primitive parameters of the model describing the regime-switching economy, market investment opportunities and investor's preferences. They solve a system of ordinary differential equations with boundary conditions equal to zero when $\tau = 0$ (*i.e.*, at maturity). We also impose the terminal condition $H(p_t, \pi_t, 0) = 1$ due to the presence of bequest. This solution is valid for the two cases with and without consumption, even though the coefficients will be different for each case. We develop the approximate Bellman equation in Appendix B by using the approximate function $H(.)$ in equation (24) for the value function and its derivatives.

Once the coefficients are obtained, the optimal portfolio strategy can be computed with the following formula:

$$\begin{aligned}
\alpha_t &= \frac{1}{\gamma}\left(\mathbf{D}_{\mathbf{s},\mathbf{t}}\pi_t - r\mathbf{1}\right)^{\mathbf{T}}\left[\left(\mathbf{V}\pi_t\right)\left(\mathbf{V}\pi_t\right)^{\mathbf{T}}\right]^{-1} \\
&+ \frac{1}{\gamma}\left[A_p(\tau) + 2B_p(\tau)p_t + \sum_{i=1}^{m} C_{pi}(\tau)\pi_{i,t}\right]\overline{\sigma}_{\mathbf{p}}\left(\mathbf{V}\pi_t\right)^{-1} \\
&+ \frac{1}{\gamma}\sum_{i=1}^{m}\left[A_i(\tau) + 2B_i(\tau)\pi_{i,t} + C_{pi}(\tau)p_t + \sum_{j\neq i} C_{ij}(\tau)\pi_{j,t}\right]\overline{\sigma}_{\mathbf{i},\pi}\left(\mathbf{V}\pi_t\right)^{-1}.
\end{aligned} \tag{25}$$

Note that all terms (myopic allocation and hedging demands) are proportional to the reciprocal of the risk aversion parameter γ. Our model implies that the Sharpe ratio is calculated using the risk premia and the volatility matrices as weighted averages of their respective values in the regimes, with the filtered probabilities as weights.

The optimal strategy for consumption is given by equation (23). Therefore optimal consumption is affected by the state variables through the function H and by ψ, the elasticity of intertemporal substitution. The portfolio allocation also depends on this parameter through the coefficients since all equations depend on ψ.

5 Assessing the Accuracy of the Approximation Method with a Four-Regime Model and Four Asset Classes

In this section we assess the accuracy of the approximation method described in the previous sections. We reproduce the model of Guidolin and Timmermann (2007) that is particularly suited to our purpose since it involves a fair degree of complexity in terms of number of regimes and asset classes. A small investor has to allocate his or her wealth between three risky asset classes, a large cap index, a small cap index and a long-term bond index, and a 30-day Treasury bill assumed to be risk-free. The asset returns evolve in four economic states, crash, slow growth, bull and recovery. They solve the strategic allocation problem with Monte Carlo methods. This will provide a benchmark solution to assess the accuracy of our approximation. While Guidolin and Timmermann (2007) estimate optimally the number of states over their sample period, we will assume that the same economic states prevail over our extended period. Since our method is developed in a continuous-time framework and the estimation has obviously to be conducted in discrete time, we refer below to an online appendix for our estimation procedure. In the following sections, we summarize our estimation results for this model and evaluate the accuracy of our approximate solution to the strategic allocation problem. We then assess the economic importance of multiple regimes and recompute the optimal portfolio decisions when the dividend yield is added as a predictor of the regimes and consumption is considered alongside the portfolio problem of the investor.

5.1 Estimation of the model parameters

To estimate the parameters, we need to write the model in discrete time. In a discrete-time strategy, the investor makes her portfolio and consumption decisions based on current information and waits until the next period to make new decisions with the newly acquired information. We set the rebalancing period equal to one month to follow the portfolio literature on discrete-time regime switching models (*e.g.*, Guidolin and Timmermann (2005, 2007, 2008a), Ang and Timmermann (2012), Ang and Bekaert (2002, 2004)).

For the large and small cap indices, we select the highest and lowest quintiles based on the capitalization value of all NYSE, AMEX, and NASDAQ firms. For the long-term bond index, we select the CRSP portfolio of 10-year Treasury bonds. The risk-free asset is the 3-month Treasury bill secondary market rate.[3] We estimate the model over the period from January 1953 to July 2018. We refer to the online appendix for a detailed description of the estimation procedure.

The parameter values are shown in Table 1. In Panel A, we show the means and correlations of the three risky asset classes assuming a single regime in the economy. The highest returns are obtained for the small caps index and the highest correlation is achieved between the large-cap index and the small-cap one. We also report the optimal portfolio shares corresponding to this single regime. We observe a slightly higher share for the large stocks (33%) than for the small stocks (28%), while bonds represent 47% of the portfolio and the balance (-8%) corresponds to the borrowing on the risk-free market.

[3] The data was extracted from the following databases: for the stock indices, Kenneth French's website; for the long-term bond index, CRSP through WRDS; and for the riskless rate, the Board of Governors of the Federal Reserve System.

In Panel B, we report the results for the four regimes. The bull-regime is good for the stocks and bad for the bonds, while the crash means are negative for the stocks and high for the bonds since there is a flight to quality in a depressed state. The slow growth affects mainly the mean returns of the small caps and the recovery boosts all asset classes with means greater than in all other states. In terms of average probabilities (see Table 2) the slow growth regime is the most frequent (around 40% of the time) followed by the bull state (around 31%). In terms of average duration (measured by $1/(1-p_{ii})$, where p_{ii} is the transition probability of staying in state i, the bull state lasts much longer than any other state, while the duration of the two extreme regimes is only of a few months.

From the crash regime, the economy transitions to the recovery state, while from the recovery state it shifts to the slow growth state. The correlations/volatilities matrix indicates that the highest volatilities are in the worst state and the low ones in the bull state. The correlation between the large and the small stocks is close to what it was for the single regime (around 73%) except in the recovery period. The correlation between bonds and small caps is negative in all regimes.

The estimation procedure provides filtered probabilities of being in each regime at time $t+1$ for each point in time t for the whole estimation period (see the graphs of these probabilities in the online appendix). These probabilities are conditional on the information available to investors when they make their portfolio and consumption decisions. We will use these conditional probabilities to evaluate the returns of dynamic strategies based on our approximate solution in Section 5.3.2.

We still have to estimate σ_π, the matrix introduced in the filtering problem in equation (6). To estimate σ_π in a discrete time setup, we first create an $(n+1)$ x m matrix (denoted by $\mathbf{D_t}$) consisting of $\mathbf{D_{s,t}}$ in the first n rows and of $\mathbf{D_{p,t}}$ over the last row. Then we calculate the monthly time-series for the drifts ($\mathbf{D_t}\pi_\mathbf{t}$) and volatility matrix ($\mathbf{V}\pi_\mathbf{t}$).[4] Then, following equations (3) and (9), we can estimate the discretized process:

$$\Delta\mathbf{Z}^*_\mathbf{t} = \sigma_\mathbf{t}{}^{-\mathbf{1}}\left(\mathbf{LR_t} - \mathbf{D_t}\pi_\mathbf{t}\right), \tag{26}$$

where $\mathbf{LR_t}$ is the $(n+1)$ x 1 vector of the log returns for the n assets together with the predictor change observed at time t (see the online appendix for more details). The matrix $\sigma_\mathbf{t}$ is the

[4] We leave out the first year of data to avoid the effect of the starting probabilities on the filtered probabilities. After a year this effect is irrelevant.

concatenation of $(\mathbf{V}\pi_\mathbf{t})$ with $\sigma_\mathbf{p}$, thus creating a square matrix of order $n+1$ (the n first elements of the last column are filled up with zeros). We use equation (6) to write the discretized process $(\Delta\pi_\mathbf{t} - \mu_{\pi,\mathbf{t}}) = \sigma_\pi \Delta\mathbf{Z}^*_\mathbf{t}$ and store the left-hand time-series in an $m \times T$ matrix denoted by $(\Delta\pi - \mu_\pi)$. We also store all increments $\Delta\mathbf{Z}^*_\mathbf{t}$ in an $(n+1) \times T$ matrix denoted simply by $\Delta\mathbf{Z}^*$ (where T is the time-series length). We finally obtain the desired estimation:

$$\sigma_\pi = (\Delta\pi - \mu_\pi)\, \Delta\mathbf{Z}^{*\mathbf{T}} \left(\Delta\mathbf{Z}^* \Delta\mathbf{Z}^{*\mathbf{T}}\right)^{-1}. \tag{27}$$

5.2 Evaluating the accuracy of the approximate solution

We specialize the general model described in the previous sections to a model without predictability and consumption. To compute the optimal portfolio solution for a dynamic investor with a finite horizon under a regime-switching economy, we rely on the numerical method proposed in Guidolin and Timmermann (2007). They solve the dynamic portfolio problem of the investor numerically by using Monte Carlo methods to approximate expected utility. We describe this method in detail in Appendix C. To use this numerical solution as a benchmark for assessing the accuracy of our approximate solution, it is essential to obtain a high level of precision. This depends mainly on how small we set the steps in the discretization grids for the probabilities and the portfolio shares.[5] Guidolin and Timmermann (2007) indicate that five grid-discretization points (steps of 20%) guarantee sufficient accuracy in the calculation of optimal choices. To be on the safe side we set our grid for each element of $\widehat{\mathbf{Y}}_{\mathbf{t+1|t}}$ with 9 discretization points $\{0, 0.125, 0.250, ..., 1\}$. For the portfolio weights $\alpha_\mathbf{t}$, we applied a very fine grid with 1% steps for all assets. We set the number of simulations at 30,000, as in Guidolin and Timmermann (2007). With this level of precision, the computing time of the numerical solution for a portfolio was more than two days in average for a particular regime (we parallelized the simulation in 12 physical CPU cores to complete it in approximate 52 processing hours for the five cases). In contrast, computing the shares of the four-asset portfolio with the approximate solution takes seconds. It took 5 to 7 seconds to compute each line of Table 3, so the whole table was produced in 2 to 3 minutes. If we find that our approximation is accurate, this large difference in computing time will be crucial when adding assets or predictors or implementing

[5] Adding a predictor would have created the need for a third grid in the numerical solution method. The complexity of considering 4 regimes and three risky assets poses already computational challenges. Similarly we did not include consumption in the problem since the problem without consumption ($\psi = \infty$) captures all the approximations we need to assess. Campani and Garcia (2019) show that the approximation needed when $\psi \neq 1$ in the Bellman equation is accurate.

investment strategies.

We report the asset shares in Table 3 under the column Optimal Shares when each of the initial states is supposed to be known and when it is unknown, in which case ergodic probabilities are used to weigh the states. The shares are estimated for different horizons, from 1 to 120 months. Note that we allow for short sales, which results in extreme positions when the initial regime is supposed to be known. In the crash state, the investor heavily sells short large stocks (less volatile than small stocks) to invest in safer assets (bonds and risk-free asset). In the slow growth regime, the investor loads on large stocks and sells short small caps and bonds, since bonds and small caps generate lower expected returns. In the bull regime, long-term bonds are hugely sold short to mainly invest in small stocks and the risk-free asset. In this regime, as in the previous one, bonds offer a lower expected return (0.14%) than the risk-free asset (0.36%). The investor borrows heavily in the recovery regime to invest mainly in small stocks. Even if the shares are unrealistic in these known-regime cases, their main purpose is to evaluate the accuracy of the proposed approximation. When the current regime is unknown, we assume that the investor uses the ergodic (long-run) probabilities of each regime. Given the uncertainty about the current regime, long-term bonds represent an important fraction of the portfolio, leverage is reduced, and overall the portfolios shares become a lot less extreme than in the known-regime cases.

Given the optimal solution obtained by simulation, we proceed to assess the accuracy of the proposed approximate analytical solution derived in Section 4.2. Our main measure of accuracy will be the total wealth equivalent utility loss created by using the approximate solution instead of the optimal solution obtained by simulation. We calculate this loss by considering two identical investors, one that follows the optimal strategy, the other the approximate strategy. The difference is that they do not invest the same initial amount. The investor who follows the approximate strategy begins with a wealth equal to \$100. Given the two investors' strategies, we match their value functions (*i.e.*, utilities), so we can calculate the actual initial wealth the optimal investor needs to start with. Given that his strategy is optimal the wealth amount will be less than \$100. The initial wealth difference will be the percentage of the total wealth equivalent utility loss due to the suboptimal strategy. If this loss is negligible, we will conclude that the approximate analytical solution provides accurate strategies.

In Table 3, the last column reports the wealth equivalent loss due to using our approximate solution instead of the optimal strategy obtained through simulation. We compute it for all horizons and for all cases of known and unknown initial regimes. For the known-regime cases, the maximum total loss is 0.02% in the most leveraged regime. When the investor wants to account for a multi-regime economy but has no access to the filtered probabilities, the maximum total loss is 0.01% when using ergodic probabilities to weigh the regimes. These figures support the reliability of the analytical approximate solution proposed in this paper.

We complement this wealth loss criterion by showing the portfolio shares computed with the approximate analytical formulas. As we simulate the optimal solution on a grid of 1% steps, for all cases, we can see that the approximate shares are very close to the optimal ones.[6]

Overall we can conclude that the approximate solution can be computed fast and offers a very good level of accuracy. One may wonder whether the approximation loses accuracy under certain parameter values. The approximation we propose is based on a second-order approximation that was used by Campbell and Viceira (1999) in discrete time for an infinite horizon. Campbell and Koo (1997) study the accuracy of the approximate solution with respect to a numerical solution and conclude that high risk aversion and a high return variability could decrease the accuracy of the log consumption-wealth ratio. Introducing regimes may be a way to capture a highly volatile continuous-time stochastic volatility process, but the precise effect on the accuracy of the approximation warrants further study.

As we will see in the next section the speed and accuracy of the approximation method become particularly important in out-of-sample dynamic portfolio allocations where shares have to be computed each month given the estimated conditional probabilities of the regimes.

5.3 The economic importance of regimes

In this section we first evaluate whether hedging demands due to regime changes are important by considering a myopic investor who ignores these changes and measuring his wealth equivalent loss with respect to an investor who follows an optimal strategy. A second way to gauge the importance of the regimes is to conduct an out-of-sample exercise to measure the extra returns generated by considering regimes with respect to a single-regime economy.

[6] This precision proved to be important. In a previous version we set the steps at 5% and the computed shares were further away from our approximate solution. We thank a referee for suggesting to use a finer grid.

5.3.1 Are hedging demands important?

In Table 3 we saw that the shares invested in each asset class varied very little with the horizon, which implies that hedging demands are relatively small. To measure their economic importance, we compute the wealth equivalent loss that an investor incurs when he follows a myopic portfolio strategy (*i.e.*, ignores hedging demands). We use the Bellman equation (30) corresponding to the portfolio myopic strategy given by:

$$\alpha_{\mathbf{t}} = \frac{1}{\gamma}\left(\mathbf{D_{s,t}}\pi_{\mathbf{t}} - \mathbf{r1}\right)^{\mathbf{T}}\left[\left(\mathbf{V}\pi_{\mathbf{t}}\right)\left(\mathbf{V}\pi_{\mathbf{t}}\right)^{\mathbf{T}}\right]^{-\mathbf{1}}. \tag{28}$$

We follow the same steps as in the evaluation of the wealth loss due to the approximate solution with respect to the optimal one. The approximate solution is given by the same value function as before, but with different coefficients. These coefficients solve a similar but simpler system of equations, which we derive in Section 5 of the online appendix where we also present that the non-predictability case is nested and obtained when $\mathbb{B}$ is a zero column matrix.

To find the total wealth equivalent loss, we equate the optimal value function (in which initial wealth is unknown) to the myopic sub-optimal one (in which wealth is standardized to 100) to solve for the unknown initial wealth of the optimal strategy. The difference with 100 gives the total percentage wealth equivalent loss due to the myopic strategy.

In Panel A of Table 4, we report the wealth equivalent loss for the myopic strategy starting from each of the four regimes and from the ergodic-probability weighted states, for horizons of 1 month, 10 years and 40 years. For all horizons, the maximum total loss remains very small (with a maximum of 0.59%). These results are known in the literature (see Ang and Bekaert (2002) and Guidolin and Timmermann (2007)). A likely reason is that in all these papers, including ours, the price of risk of switching regimes is assumed to be zero. Several studies have shown that the part of the risk premia due to this risk could be sizable: for a recent paper, see Baele, Bekaert, Inghelbrecht and Wei (2020); for a precursor article, see Mayfield (2004); and in continuous time, see Elliott and Siu (2011).

5.3.2 Out-of-sample performance

We conduct an illustrative out-of-sample exercise to compare the performance of a four-regime model with respect to a single-regime model and a $1/N$ portfolio. The first is a natural benchmark while the second beats a number of models based on optimal rules (see De Miguel et al. (2009)).

In our setting, for the regime models, a risk-averse investor determines each month the optimal allocation for a fixed 10-year horizon given his available information and evaluates the respective returns over the next month. For the $1/N$ model, the investor simply rebalances his portfolio to maintain the $1/N$ allocation. In a real-time application, investors can only use the information available at time t to compute the allocation for time $t+1$. Therefore, mean and variance-covariance parameters have to be re-estimated each month. For simplicity, we depart from this by giving investors previous knowledge of the full-sample parameters. This of course will produce an optimistic bias in the performance of the regime models and more so for the multi-regime model.[7] However, we want to illustrate that our approximate method for computing portfolio shares is particularly suited to time-varying investment opportunities. Given the value of the mean and variance-covariance parameters, we do incorporate each month the new information in the computation of the conditional regime probabilities and provide a solution for the optimal portfolio much faster than the currently available numerical methods.

In Panel A of Table 5, we report the first four moments of the unconditional one-month return distributions for the three portfolio allocation strategies and three samples starting in 1960, 1980 and 2000 and ending in 2018. While the means of the four-regime strategy are roughly double the size of the means for the single-regime or equal-weight strategies, their variance and excess kurtosis, features that investors do not like, are much bigger. The differences in skewness between the strategies are less extreme, except for the sub-sample starting in 2000, where it is strongly positive for the multi-regime strategy, a good thing for investors. These statistics tell us that the distribution of the one-month out-of-sample returns is definitely non-normal for the multi-regime strategy and that the two other strategies show some small departure from normality. To compare and rank strategies across different samples, we compute the annualized out-of-sample certainty-equivalent return (CER) defined as:

$$\mathrm{CER}_i(\gamma, T) \equiv \frac{12}{T}\left\{\frac{1}{W_t}\left[\frac{1}{K-T}\sum_{\tau=1}^{K-T}[W_{\tau+T}(\hat{\alpha}_{i,t}(\gamma,T))]^{1-\gamma}\right]^{\frac{1}{1-\gamma}} - 1\right\}, \tag{29}$$

where T is the horizon (10 years), K is the number of out-of-sample returns, γ is the relative risk aversion parameter of a power utility function (set to 5), $\hat{\alpha}_{i,t}$ are the proportions of the wealth invested in asset i, and W_t is the initial wealth (set to 1). We provide the derivation of this

[7] Indeed we minimize estimation errors which are more costly for the high number of parameters in the four-regime model.

expression in the online appendix (see also Fugazza, Guidolin and Nicodano (2015)).

The values of the annualized CER differences between the four-regime and the two benchmark strategies for the three samples are reported in Panel B of Table 5. For the first two samples, the differences are negative (between -7% and -10%), meaning that the multi-regime strategy is dominated by the single-regime and the 1/N strategies. The large variance and kurtosis penalize strongly the four-regime strategy. For the last sample since 2000, the multi-regime allocation provides a higher CER than the two other strategies, with a positive difference of around 7% in both cases. Given the high turnover shown in Table 5 it is expected that the introduction of transaction costs will erase some of the extra returns provided by the strategy.

We also report below these figures a 95% bootstrap confidence interval. Given the non-normalities observed in the out-of-sample returns, we use the bias-corrected and accelerated percentile method.[8] For all pairs and samples, zero is outside the bounds of the confidence intervals, which means that the CER differences are statistically different from zero.

In the online appendix, we provide another example where we compute the approximate optimal shares for a five-asset portfolio allocation between three value portfolios, a bond index and a risk-free asset, also in presence of four regimes. The out-of-sample returns exhibit less variance and kurtosis and the multi-regime strategy always produces a higher CER than the two other strategies but the differences are never significantly different from zero.

5.4 Including predictability

In this section, we include predictability as formulated in the general model of Section 3. As in Guidolin and Timmermann (2007), Campbell and Viceira (1999) and Campbell et al. (2003), among others, we choose the dividend yield as a predictor of asset returns and use the logarithm of the S&P 500 dividend yield as a measure. We report the parameter estimates of the asset returns and the predictor dynamics together with the transition probabilities in Table 6. Overall, the information provided by the predictor allows investors to make more precise inferences about the regimes, but does not affect in an important way the magnitude of the means, volatilities and correlations of asset returns compared with the no-predictability case. However, we notice that the regime probabilities have changed slightly, with the crash and the slow growth being less prevalent and leaving more weight for the good states. At the bottom of

[8] We draw 10,000 samples with replacement of the same size as the respective data samples (from 1960, 1980 and 2000 till 2018). Given the low values of the first-order autocorrelation coefficient reported in Panel A of Table 5, we did not judge necessary to use a block bootstrap.

Table 6, we report the estimated speed of mean reversion of the Ornstein-Uhlenbeck process for the dividend yield. Its value of 0.0163 at the monthly frequency implies a half-life of about 3.5 years.

5.4.1 Dynamic strategies with predictability and sensitivity to the predictor

We plot the various strategies for the four asset classes and all monthly horizons up to ten years in Figure 1. Let us start with the case where the value of the predictor is set at its mean value. Compared with the case without predictability, we see at short horizon that the overall investment share in stocks (large and small) increases from 85.6% to 93.4%, at the expense mainly of bonds. However, at a horizon of 10 years, the stock share decreases to 82.3%. This is mainly due to the mean reversion in the dividend yield. We can see that the effect is strongly present in the large stocks since the dividend yield of the S&P 500 is mostly related to the large stocks. The share increases up to the horizon of 3 to 4 years, which is consistent with the persistence of the dividend yield process, and then decreases to about 27% of the portfolio. On the contrary, the share of the small stocks increases by 5% up to the 10-year horizon since their mean returns are higher than for the large stocks.

The level of the predictor with which we compute the shares is of course key to the optimal portfolio strategies. Our approximate formulas allow us to compute easily the shares for all horizons when we set the predictor at different values. Here we add or subtract two standard deviations to the mean value of the predictor. Overall, the share invested in stocks is higher for all horizons when the dividend yield is two standard deviations above its mean (from 106% to 103%) but the patterns for large and small stocks are quite different. The share is small stocks is much higher and increases from 73% to 82% of the invested portfolio. This is mainly due to the higher predictor coefficient in the drift of the small stocks (vector $\mathbb{B}$ in Table 6). For the large stocks, we observe the same pattern as before but at a lower level. For bonds, a higher predictor level makes them more attractive with respect to the Treasury bills (again through the effect on their drift) and the investor borrows more to take advantage of the wedge in returns. The allocations are reversed when the dividend yield is two standard deviations below its mean. Large stocks take preminence over small stocks and the investors hedges his portfolio by selling bonds and buying Treasury bills.

5.4.2 Predictability and hedging demands

Using the same methodology as before, we calculate the expected wealth equivalent loss that a myopic investor suffers in the beginning of the investment period by ignoring hedging demands. Results are reported in Panel B of Table 4 . We observe that when the predictor is equal to its sample average, above or below its mean, the total loss is always less than 0.04% for a one-year horizon, which is negligible. The losses increase significantly for a 10-year horizon for regime 3, regime 4 and the ergodic probability case. For a 40-year strategy, hedging demands become large, as the cost of being myopic represents losses in the order of 90% when the regime is unknown.

5.5 Including consumption

In this section, we determine how the presence of multiple regimes affects the consumption pattern of investors. We report in Figure 2 how optimal consumption evolves with and without predictability, for realistic values of ψ and of the predictor.[9] We also analyzed the impact of including consumption on the portfolio strategy and concluded that it was small with and without predictability. Therefore, all previous results hold also with intermediate consumption.

In the upper panel, we show the consumption policy as a function of the horizon for each regime as well as when the investor uses ergodic or equal probabilities (the predictor is set at its historical mean, ψ is fixed at 0.75 and γ remains, as always in this paper, equal to 5). We can observe that consumption is roughly equal in both dynamic problems with and without predictability. Moreover, and somehow surprisingly, consumption policy barely changes with the regime.

The lower plots present the real-time exercise of how a 12-month horizon investor would consume with the actual filtered probabilities and predictor values from July 1988 to July 2018. In these graphs, we show five different investors: two willing to postpone consumption ($\psi = 1.25$ and $\psi = 1.5$), two others that hardly substitute consumption today for tomorrow ($\psi = 0.5$ and $\psi = 0.75$) and the myopic (in terms of consumption) investor ($\psi = 1$). These two plots are important because they use a broad real-time range for current probabilities and for the predictor. They confirm that consumption-to-wealth ratio variations with the predictor and with the regimes are very limited.

[9] For a similar problem without regimes Campani and Garcia (2019) analyze in detail the sensitivity of consumption and portfolio choices to the value of both preference parameters γ and ψ.

6 Adding Assets or Predictors and Computation Time of the Approximate Method

We have shown in the previous sections that our approximate analytical approach is much faster than the numerical methods currently used for portfolio allocations with regime-switching asset returns, without any significant loss in accuracy. One question remains. Can we keep or increase the speed advantage relative to the numerical methods when we increase the number of assets or predictors? The answer is unequivocally yes in terms of computation time. Indeed, the approximate solution for the case with no predictability is obtained by solving a system of 15 differential equations, given four regimes, to compute the coefficients A_0 (the independent term), A_is corresponding to the probabilities $\pi_{i,t}$, B_is corresponding to the squares of the probabilities $\pi_{i,t}^2$, and finally $C_{ij}, i < j$ corresponding to the cross-products of the probabilities $\pi_{i,t}\pi_{j,t}$. It is clear that these computations are practically independent of the number of assets. To solve the problem with a predictor we need to add 6 equations for the coefficients A_p, B_p and $C_{p,i}$ corresponding to p_t, p_t^2 and $p_t\pi_{i,t}$ (see Appendix B and Section 4 of the online appendix for details). Adding another predictor will create another set of six equations. The solution for these larger systems will still be computed in a few minutes. In numerical methods, the computation cost increases considerably with the number of assets or predictors since an additional grid has to be added for each new variable considered in the system.

The limitation to increasing the number of assets, in a multi-regime model in particular, is the increase of the number of parameters that have to be estimated. Solutions may be found in using indices that capture regimes in all asset returns. For example, using the VIX to identify changes in economic or financial regimes may offer a way to increase the number of asset classes without increasing the estimation burden. Other indices measuring tail risk or funding liquidity risk could be used in other applications. Our method will then offer a real computation advantage compared with existing numerical methods.

7 Conclusion

This paper derives an approximate analytical solution to a dynamic portfolio and consumption problem under stochastic differential utility in a multi-regime economy. We have shown that the method is accurate in a comparison with a numerical solution of the four-asset, four-regime model of Guidolin and Timmermann (2007), which features a portfolio allocation between large and small stocks, long-term government bonds and a risk-free asset. The main advantage of the method is its speed of computation of the portfolio shares. Moreover, without predictability, the speed of computation will be the same for a larger number of assets. Adding a predictor does not increase the computation time significantly. The method is therefore potentially very useful for real-time portfolio management. However, at least three problems remain. With a large number of assets, the parameter estimation problem for the dynamics of returns in the various regimes becomes overwhelming. Another issue is the large turnover and the increase in transaction costs that optimal portfolio allocation in the presence of regimes entails. A third problem is that, for some funds, short-sale constraints have to be imposed. Designing solutions to address these three problems is left for future research.

References

Ang, Andrew and Bekaert, Geert (2002). “International Asset Allocation with Regime Shifts”, *Review of Financial Studies*, Vol.15, No. 4, 1137-1187.

Ang, Andrew and Bekaert, Geert (2004). “How Regimes Affect Asset Allocation”, *Financial Analysts Journal*, Vol. 60, No. 2, 86-99.

Ang, Andrew and Timmermann, Allan (2012). “Regime Changes and Financial Markets”, *Annual Review of Financial Economics*, Vol. 4, 313-337.

Baelel, Lieven; Bekaert, Geert; Inghelbrecht, Koen; and Wei, Min (2020). “Flights to Safety”, *Review of Financial Studies*, Vol.33, No. 2, 610-643.

Berrada, Tony; Detemple, Jérôme; and Rindisbacher, Marcel (2018). “Asset Pricing with Beliefs-dependent Risk Aversion and Learning”, *Journal of Financial Economics*, Vol. 128, No. 3, 504-534.

Campani, Carlos H. and Garcia, René (2019). “Approximate Analytical Solutions for Consumption and Portfolio Decisions under Recursive Utility and Finite Horizon”, *The North American Journal of Economics and Finance*, Vol. 48, 364-384.

Campbell, John Y. and Viceira, Luis M. (1999). “Consumption and Portfolio Decisions When

Expected Returns Are Time Varying", *Quarterly Journal of Economics*, Vol. 114, No. 2, 433-495.

Campbell, John Y. and Koo, Hyeng Keun (1997). "A Comparison of Numerical and Analytic Approximate Solutions to an Intertemporal Consumption Choice Problem", *Journal of Economic Dynamics and Control*, Vol. 21, 273-295.

Campbell, John Y.; Chan, Yeung L.; and Viceira, Luis M. (2003). "A Multivariate Model of Strategic Asset Allocation", *Journal of Financial Economics*, Vol. 67, No. 1, 41-80.

Campbell, John Y.; Chacko, George; Rodriguez, Jorge; and Viceira, Luis M. (2004). "Strategic Asset Allocation in a Continuous-Time VAR Model", *Journal of Economic Dynamics and Control*, Vol. 28, No. 11, 2195-2214.

Çanakoğlu, Ethem and Özekici, Süleyman (2012). "HATA Frontiers of Optimal Portfolios in Stochastic Markets", *European Journal of Operational Research*, Vol. 221, No. 1, 129-137.

David, Alexander (1997). "Fluctuating Confidence in Stock Markets: Implications for Returns and Volatility", *Journal of Financial and Quantitative Analysis*, Vol. 32, No. 4, 427-462.

David, Alexander and Veronesi, Pietro (2013). "What Ties Return Volatilities to Price Valuations and Fundamentals?", *Journal of Political Economy*, Vol. 121, No. 4, 682-746.

De Miguel, Victor; Garlappi, Lorenzo; and Uppal, Raman (2009). "Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?", *The Review of Financial Studies*, Vol. 22, No. 5, 1915-1953.

Duffie, Darrell and Epstein, Larry G. (1992). "Stochastic Differential Utility", *Econometrica*, Vol. 60, No. 2, 353-394.

Elliott, Robert. J. and, Siu, Tak Kuen (2011). "Pricing and Hedging Contingent Claims with Regime Switching Risk.", *Communications in Mathematical Sciences*, Vol. 9, 477–498.

Fugazza, Carolina; Guidolin, Massimo and Nicodano, Giovanna (2015). "Equally Weighted vs. Long-Run Optimal Portfolios.", *European Financial Management*, Vol. 21, No. 4, 742-789.

Graflund, Andreas and Nilsson, Birger (2003). "Dynamic Portfolio Selection: The Relevance of Switching Regimes and Investment Horizon", *European Financial Management*, Vol. 9, No. 2, 179-200.

Guidolin, Massimo and Hyde, Stuart (2012). "Can VAR Models Capture Regime Shifts in Asset Returns? A Long-Horizon Strategic Asset Allocation Perspective", *Journal of Banking and Finance*, Vol. 36, No. 3, 695-716.

Guidolin, Massimo and Timmermann, Allan (2005). "Strategic Asset Allocation and Consumption Decisions under Multivariate Regime Switching", Working paper.

Guidolin, Massimo and Timmermann, Allan (2006). "Optimal Portfolio Choice under Regime Switching, Skew and Kurtosis Preferences", Working paper.

Guidolin, Massimo and Timmermann, Allan (2007). "Asset Allocation under Multivariate Regime Switching", *Journal of Economic Dynamics and Control*, Vol. 31, No. 11, 3503-3544.

Guidolin, Massimo and Timmermann, Allan (2008a). "International Asset Allocation under Regime Switching, Skew and Kurtosis Preferences", *Review of Financial Studies*, Vol. 21, No. 2, 889-935.

Guidolin, Massimo and Timmermann, Allan (2008b). "Size and Value Anomalies under Regime Shifts", *Journal of Financial Econometrics*, Vol. 6, No. 1, 1-48.

Hamilton, James D. (1989). "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica*, Vol. 57, No. 2, 357-384.

Haussmann, Ulrich and Sass, Jörn (2004). "Optimal Terminal Wealth under Partial Information for HMM Stock Returns", *Contemporay Mathematics*, Vol. 351, 171-185.

Honda, Toshiki (2003). "Optimal Portfolio Choice for Unobservable and Regime-Switching Mean Returns", *Journal of Economic Dynamics and Control*, Vol. 28, No. 1, 45-78.

Kraft, Holger; Seiferling, Thomas; and Seifried, Frank T. (2017). "Optimal Consumption and Investment with Epstein-Zin Recursive Utility", *Finance and Stochastics*, Vol. 21, No. 1, 187-226.

Krishnamurthy, Vikram; Leoff, Elisabeth; and Sass, Jörn (2018). "Filterbased Stochastic Volatility in Continuous-time Hidden Markov Models", *Econometrics and Statistics*, Vol. 6, 1-21.

Lim, Andrew E.B. and Watewai, Thaisiri (2012). "Optimal Investment and Consumption when Regime Transition Cause Price Shocks", *Insurance: Mathematics and Economics*, Vol. 51, No. 3, 551-566.

Liu, Hening (2011). “Dynamic Portfolio Choice under Ambiguity and Regime Switching Mean Returns”, *Journal of Economic Dynamics and Control*, Vol. 35, No. 4, 623-640.

López, Oscar and Serrano, Rafael (2015). “Martingale Approach to Optimal Portfolio-Consumption Problems in Markov-Modulated Pure-Jump Models”, *Stochastic Models*, Vol. 31, No. 2, 261-291.

Mayfield, E. Scott (2004). “Estimating the Market Risk Premium”, *Journal of Financial Economics*, Vol. 73, No. 3, 465–496.

Sotomayor, Luz R. and Cadenillas, Abel (2009). “Explicit Solutions of Consumption-Investment Problems in Financial Markets with Regime Switchings”, *Mathematical Finance*, Vol. 19, No. 2, 251-279.

Tu, Jun (2010). “Is Regime Switching in Stock Returns Important in Portfolio Decisions?”, *Management Science*, Vol. 56, No. 7, 1198-1215.

Veronesi, Pietro (2000). “How Does Information Quality Affect Stock Returns?”, *Journal of Finance*, Vol. 55, No. 2, 807-837.

Wonham, Walter M. (1964). “Some Applications of Stochastic Differential Equations to Optimal Nonlinear Filtering”, *SIAM Journal of Control*, Vol. 2, No. 3, 347-368.

Xing, Hao (2017). “Consumption-Investment Optimization with Epstein-Zin Utility in Incomplete Markets”, *Finance and Stochastics*, Vol. 21, No. 1, 227-262.

Yin, George and Zhou, Xun Yu (2003). “Markowitz’s Mean-Variance Portfolio Selection with Regime Switching: A Continuous-Time Model”, *SIAM Journal on Control and Optimization*, Vol. 42, No.4, 1466-1482.

Appendices

A The Bellman Equation for the Portfolio Problem Without Consumption

With no intermediate consumption, the Bellmann equation simplifies to the following expression:

$$0 = \sup_{\{\alpha_t\}} \begin{bmatrix} -\beta J_t + \frac{\partial J_t}{\partial t} + J_w W_t \left[r + \alpha_t \left(\mathbf{D_{s,t}} \pi_t - r\mathbf{1} \right) \right] \\ +\frac{1}{2} J_{ww} W_t^2 \alpha_t \left(\mathbf{V}\pi_t \right) \left(\mathbf{V}\pi_t \right)^{\mathbf{T}} \alpha_t^{\mathbf{T}} + J_p \mathbf{D_{p,t}} \pi_t + \frac{1}{2} J_{pp} \sigma_{\mathbf{p}} \sigma_{\mathbf{p}}^{\mathbf{T}} \\ + \sum_{i=1}^{m} J_{\pi_i} \sum_{j=1}^{m} \lambda_{ji} \pi_{j,t} + \frac{1}{2} \sum_{i=1}^{m} J_{\pi_i \pi_i} \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{i},\pi}^{\mathbf{T}} + J_{wp} W_t \alpha_t \left(\mathbf{V}\pi_t \right) \overline{\sigma}_{\mathbf{p}}^{\mathbf{T}} \\ + \sum_{i=1}^{m} J_{w\pi_i} W_t \alpha_t \left(\mathbf{V}\pi_t \right) \overline{\sigma}_{\mathbf{i},\pi}^{\mathbf{T}} + \sum_{i=1}^{m} J_{p\pi_i} \sigma_{\mathbf{p}} \sigma_{\mathbf{i},\pi}^{\mathbf{T}} + \sum_{i<j} J_{\pi_i \pi_j} \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{j},\pi}^{\mathbf{T}} \end{bmatrix}. \tag{30}$$

Substituting the optimal expression for the portfolio weight, as given by equation (20), the final Bellman equation (without consumption) is given by:

$$\begin{aligned} 0 \;=\; & -\beta V_t + \frac{\partial V_t}{\partial t} + V_w W_t r + V_p \mathbf{D_{p,t}} \pi_t + \frac{1}{2} V_{pp} \sigma_{\mathbf{p}} \sigma_{\mathbf{p}}^{\mathbf{T}} \\ & + \sum_{i=1}^{m} V_{p\pi_i} \sigma_{\mathbf{p}} \sigma_{\mathbf{i},\pi}^{\mathbf{T}} + \sum_{i,j=1}^{m} V_{\pi_i} \lambda_{ji} \pi_{j,t} + \frac{1}{2} \sum_{i,j=1}^{m} \left(V_{\pi_i \pi_j} \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{j},\pi}^{\mathbf{T}} - \frac{V_{w\pi_i} V_{w\pi_j}}{V_{ww}} \overline{\sigma}_{\mathbf{i},\pi} \overline{\sigma}_{\mathbf{j},\pi}^{\mathbf{T}} \right) \\ & - \frac{1}{2} \frac{V_w^2}{V_{ww}} \left(\mathbf{D_{s,t}} \pi_t - r\mathbf{1} \right)^{\mathbf{T}} \left[\left(\mathbf{V}\pi_t \right) \left(\mathbf{V}\pi_t \right)^{\mathbf{T}} \right]^{-\mathbf{1}} \left(\mathbf{D_{s,t}} \pi_t - r\mathbf{1} \right) - \frac{1}{2} \frac{V_{wp}^2}{V_{ww}} \overline{\sigma}_{\mathbf{p}} \overline{\sigma}_{\mathbf{p}}^{\mathbf{T}} \\ & - \frac{V_w V_{wp}}{V_{ww}} \overline{\sigma}_{\mathbf{p}} \left(\mathbf{V}\pi_t \right)^{-\mathbf{1}} \left(\mathbf{D_{s,t}} \pi_t - r\mathbf{1} \right) - \sum_{i=1}^{m} \frac{V_w V_{w\pi_i}}{V_{ww}} \overline{\sigma}_{\mathbf{i},\pi} \left(\mathbf{V}\pi_t \right)^{-\mathbf{1}} \left(\mathbf{D_{s,t}} \pi_t - r\mathbf{1} \right) \\ & - \sum_{i=1}^{m} \frac{V_{wp} V_{w\pi_i}}{V_{ww}} \overline{\sigma}_{\mathbf{p}} \overline{\sigma}_{\mathbf{i},\pi}^{\mathbf{T}}. \end{aligned} \tag{31}$$

B The Approximate Bellman Equation

To solve the Bellman equation, we need to make the volatility matrix $\mathbf{V}\pi_t$ constant. A natural way to do it is to take the unconditional expectation of Y_t, in other words to replace the transition probabilities π_t by the unconditional regime probabilities. We denote the corresponding volatility matrix by $\mathbf{V}\pi_\infty$. Also, it is convenient to work with a row vector of conditional predictor means in each regime $\bar{\mathbb{P}} = \begin{bmatrix} \bar{p}_1 & \bar{p}_2 & \dots & \bar{p}_m \end{bmatrix}$.

The Bellman equation obtained with the approximate function $H(p_t, \pi_t, \tau) = e^{g(p_t, \pi_t, \tau)}$ defined in equation (24) is as follows:

$$
\begin{aligned}
0 \;=\;& \hat{f}(g) - A_0' - A_p' p_t - \sum_{i=1}^{m} A_i' \pi_{i,t} - B_p' p_t^2 - \sum_{i=1}^{m} B_i' \pi_{i,t}^2 - \sum_{i=1}^{m} C_{pi}' p_t \pi_{i,t} \\
& - \sum_{i<j} C_{ij}' \pi_{i,t}\pi_{j,t} + (1-\gamma)\, r \;+ \left(A_p + 2B_p p_t + \sum_{i=1}^{m} C_{pi}\pi_{i,t} \right) * \kappa_p (\bar{\mathbb{P}}\pi_t - p_t) \\
& + \frac{1}{2}\left[2B_p + \left(A_p + 2B_p p_t + \sum_{i=1}^{m} C_{pi}\pi_{i,t} \right)^2 \right] \sigma_{\mathbf{p}} \sigma_{\mathbf{p}}^{\mathbf{T}} \\
& + \sum_{i=1}^{m} \left[C_{pi} + \left(A_p + 2B_p p_t + \sum_{k=1}^{m} C_{pk}\pi_{k,t} \right) \left(A_i + 2B_i \pi_{i,t} + C_{pi} p_t + \sum_{k \neq i} C_{ik}\pi_{k,t} \right) \right] \sigma_{\mathbf{p}} \sigma_{\mathbf{i},\pi}^{\mathbf{T}} \\
& + \frac{1-\gamma}{\gamma} \sum_{i=1}^{m} \left(A_p + 2B_p p_t + \sum_{k=1}^{m} C_{pk}\pi_{k,t} \right) \left(A_i + 2B_i \pi_{i,t} + C_{pi} p_t + \sum_{k \neq i} C_{ik}\pi_{k,t} \right) \bar{\sigma}_{\mathbf{p}} \bar{\sigma}_{\mathbf{i},\pi}^{\mathbf{T}} \\
& + \sum_{i,j=1}^{m} \lambda_{ji}\pi_{j,t} \left(A_i + 2B_i \pi_{i,t} + C_{pi} p_t + \sum_{k \neq i} C_{ik}\pi_{k,t} \right) + \sum_{i=1}^{m} B_i \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{i},\pi}^{\mathbf{T}} + \sum_{i<j} C_{ij} \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{j},\pi}^{\mathbf{T}} \\
& + \frac{1}{2} \sum_{i,j=1}^{m} \left(A_i + 2B_i \pi_{i,t} + C_{pi} p_t + \sum_{k \neq i} C_{ik}\pi_{k,t} \right) \left(A_j + 2B_j \pi_{j,t} + C_{pj} p_t + \sum_{k \neq j} C_{jk}\pi_{k,t} \right) \sigma_{\mathbf{i},\pi} \sigma_{\mathbf{j},\pi}^{\mathbf{T}} \\
& + \frac{1}{2}\frac{1-\gamma}{\gamma} \sum_{i,j=1}^{m} \left(A_i + 2B_i \pi_{i,t} + C_{pi} p_t + \sum_{k \neq i} C_{ik}\pi_{k,t} \right) \left(A_j + 2B_j \pi_{j,t} + C_{pj} p_t + \sum_{k \neq j} C_{jk}\pi_{k,t} \right) \bar{\sigma}_{\mathbf{i},\pi} \bar{\sigma}_{\mathbf{j},\pi}^{\mathbf{T}} \\
& + \frac{1}{2}\frac{1-\gamma}{\gamma} \left(\mathbb{A}\pi_{\mathbf{t}} + \mathbb{B} p_t - r\mathbf{1} \right)^{\mathbf{T}} \left[\left(\mathbf{V}\pi_\infty \right) \left(\mathbf{V}\pi_\infty \right)^{\mathbf{T}} \right]^{-\mathbf{1}} \left(\mathbb{A}\pi_{\mathbf{t}} + \mathbb{B} p_t - r\mathbf{1} \right) \\
& + \frac{1}{2}\frac{1-\gamma}{\gamma} \left(A_p + 2B_p p_t + \sum_{i=1}^{m} C_{pi}\pi_{i,t} \right)^2 \bar{\sigma}_{\mathbf{p}} \bar{\sigma}_{\mathbf{p}}^{\mathbf{T}} \\
& + \frac{1-\gamma}{\gamma} \left(A_p + 2B_p p_t + \sum_{i=1}^{m} C_{pi}\pi_{i,t} \right) \bar{\sigma}_{\mathbf{p}} \left(\mathbf{V}\pi_\infty \right)^{-\mathbf{1}} \left(\mathbb{A}\pi_{\mathbf{t}} + \mathbb{B} p_t - r\mathbf{1} \right) \\
& + \frac{1-\gamma}{\gamma} \sum_{i=1}^{m} \left(A_i + 2B_i \pi_{i,t} + C_{pi} p_t + \sum_{k \neq i} C_{ik}\pi_{k,t} \right) \bar{\sigma}_{\mathbf{i},\pi} \left(\mathbf{V}\pi_\infty \right)^{-\mathbf{1}} \left(\mathbb{A}\pi_{\mathbf{t}} + \mathbb{B} p_t - r\mathbf{1} \right),
\end{aligned}
\tag{32}
$$

where:

$$
\begin{aligned}
\hat{f}(g) &= -\beta && \text{with no consumption,} \\
\hat{f}(g) &= \frac{1-\gamma}{\psi-1}\left(\beta^{\psi} e^{\frac{1-\psi}{1-\gamma} g} - \beta\psi \right) && \text{when } \psi \neq 1, \text{ and} \\
\hat{f}(g) &= (1-\gamma)\,\beta \left[\ln\left(\beta e^{\frac{-1}{1-\gamma} g} \right) - 1 \right] && \text{when } \psi = 1.
\end{aligned}
$$

For the problems with no consumption or with $\psi = 1$, this equation can be solved by matching equivalent terms, such that we find a system of $\frac{m^2}{2} + \frac{5m}{2} + 3$ equations (where m represents the number of regimes) with exactly the same number of (time-varying) unknowns (the coefficients of function g). The boundary condition for all coefficients is zero when $\tau = 0$. In the online appendix to this paper, we present the system of equations which can be solved with standard software. For the general problem in which $\psi \neq 1$, this equation is non-linear, what prevents us to solve it analytically. However, using the same log-linearization approximation technique explained in Campani and Garcia (2019), we are able to solve it. Details for this general solution can also be found in the online appendix. But when $\gamma = 1$, the solution to the system (32) implies that all coefficients are constant and equal to zero, except $A_0 = -\beta\tau$. As a consequence, the optimal portfolio is myopic (no hedging demands), as expected. Another nested case is the single state model: in such a case, $\pi_{\mathbf{i}}$ is constant and therefore $\sigma_{\mathbf{i},\pi}$ is zero (no hedging demands related to regime changes).

C Details on the Monte Carlo Simulation Procedure

Starting with the dynamic problem, fix the starting point of the problem at $t = 0$, an horizon T, and a rebalancing frequency (which will be monthly). The investor's problem at any given moment t is:[10]

$$V\left(W_t, \widehat{\mathbf{Y}}_{\mathbf{t+1|t}}, t\right) = \sup_{\{\alpha_{\mathbf{t}}, \alpha_{\mathbf{t+1}}, \ldots, \alpha_{\mathbf{T-1}}\}} E_t\left[\frac{W_T^{1-\gamma}}{1-\gamma}\right], \qquad t = 0, 1, \ldots, T-1. \tag{33}$$

At date 0, the investor wants to find the optimal asset allocation $\alpha_{\mathbf{0}}$, given her horizon T, knowing that she will optimally reallocate again her resources monthly from $t = 1$ towards $t = T - 1$. The wealth constraint reads:

$$W_{t+1} = W_t\left[\left(1 - \alpha_{\mathbf{t}}^{\mathbf{T}}\mathbf{1}\right)e^r + \alpha_{\mathbf{t}}^{\mathbf{T}} e^{\overline{\mathbf{LR}}_{\mathbf{t}}}\right], \tag{34}$$

in which $e^{\overline{\mathbf{LR}}_{\mathbf{t}}}$ is a column vector where its elements are the exponentials of the elements of the vector $\mathbf{LR_t}$, defined in Section 5.1, without its last element (predictor change). With power

[10] The state variables $\widehat{\mathbf{Y}}_{\mathbf{t+1|t}}$ are the filtered probabilities in the discrete-time model and represent the best guess of the probabilities to the next period, given the information available at the current period. The equations used to compute these probabilities can be found in the online appendix.

utility, the value function can be written as:

$$V\left(W_t, \widehat{\mathbf{Y}}_{\mathbf{t+1|t}}, t\right) = \frac{W_t^{1-\gamma}}{1-\gamma} H\left(\widehat{\mathbf{Y}}_{\mathbf{t+1|t}}, t\right), \tag{35}$$

with terminal condition $H\left(\widehat{\mathbf{Y}}_{\mathbf{T+1|T}}, T\right) = 1$. Equations (33) and (35) imply the following recursion:

$$H\left(\widehat{\mathbf{Y}}_{\mathbf{t+1|t}}, t\right) = \sup_{\{\alpha_t\}} E_t\left[\left(\frac{W_{t+1}}{W_t}\right)^{1-\gamma} H\left(\widehat{\mathbf{Y}}_{\mathbf{t+2|t+1}}, t+1\right)\right] \times sign(1-\gamma), \tag{36}$$

where $sign(1-\gamma)$ is either $+1$ if $\gamma < 1$ or -1 if $\gamma > 1$.

This recursive equation is key to solve the problem using dynamic programming techniques. We start with the last investor decision, at time $t = T-1$: in this case we know that $H\left(\widehat{\mathbf{Y}}_{\mathbf{T+1|T}}, T\right) = 1$ for all $\widehat{\mathbf{Y}}_{\mathbf{T+1|T}}$ and the right-hand side of the recursion will only depend on wealth's return. The next step is to build a grid of possible values for the vector of regime probabilities ($\widehat{\mathbf{Y}}_{\mathbf{T|T-1}}$): Guidolin and Timmermann (2007) indicate in footnote 12 on page 3523 that five grid-discretization points (steps of 20%) guarantee sufficient accuracy in the calculation of optimal choices. To be on the conservative side we set our grid for each element of $\widehat{\mathbf{Y}}_{\mathbf{t+1|t}}$ with 9 discretization points $\{0, 0.125, 0.250, ..., 1\}$. For each point in the grid (the full vector $\widehat{\mathbf{Y}}_{\mathbf{t+1|t}}$), we first simulate the next period regime (using the transition probability matrix) and, based on this simulated regime, simulate the corresponding risky assets returns. To compute the right-hand side of the recursion above, we do this simulation $\mathcal{N}$ times and use the law of large numbers to approximate the expectation. We use 30,000 simulations as in Guidolin and Timmermann (2007) (see the above-referenced footnote).

To find the optimal allocation strategy, we work again with a grid for the portfolio weights α_t. We applied a very fine grid with 1% steps for all assets. Having the right-hand side of the above recursion for every $\widehat{\mathbf{Y}}_{\mathbf{T|T-1}}$ in the grid, we know $H\left(\widehat{\mathbf{Y}}_{\mathbf{T|T-1}}, T-1\right)$ for every point in the probability grid. We repeat the process backwards until we reach $t = 0$, but at this point we know the value of $\widehat{\mathbf{Y}}_{\mathbf{1|0}}$ such that the optimal portfolio strategy at time 0, given the monthly rebalancings and the horizon T, is determined. One last detail remains to be explained: not at time $t = T-1$, but at all proceeding steps, given the grid for $\widehat{\mathbf{Y}}_{\mathbf{t+1|t}}$, nothing guarantees that $\widehat{\mathbf{Y}}_{\mathbf{t+2|t+1}}$ will also be in the grid. However we need the estimated values of $H\left(\widehat{\mathbf{Y}}_{\mathbf{t+2|t+1}}, t+1\right)$ but these values were calculated only for points in the grid. We set $\widetilde{\mathbf{Y}}_{\mathbf{t+2|t+1}}$ to the closest $\widehat{\mathbf{Y}}_{\mathbf{t+2|t+1}}$ in the probability grid (using the standard Euclidean distance).

D Tables

Table 1: Parameter Estimates of the Four-Regime Model with Four Assets. We report in Panel A the parameters and the optimal weights for the single state model (for $\gamma = 5$). The optimal weights for the single-state model sum up to 100% by short selling the riskless asset. Panel B presents the monthly parameter estimates for the four-state model from January 1953 to July 2018. Both panels do not allow for predictability. The volatilities are reported on the diagonals of the correlations/volatilities matrices. The average Treasury yield over the full period was estimated at 0.36% a month.

Panel A: Single-State Model	Large Caps	Small Caps	LT Bonds	
Mean Returns	0.92%	1.15%	0.49%	
Correlations/Volatilities Matrix				
Large Caps	4.08%			
Small Caps	73.94%	5.90%		
LT Bonds	8.76%	-3.67%	2.11%	
Optimal Weights	33.39%	27.58%	46.95%	
Panel B: Four-State Model	Large Caps	Small Caps	LT Bonds	
Mean Returns				
Regime 1 (Crash)	-1.33%	-1.82%	0.71%	
Regime 2 (Slow Growth)	0.92%	0.50%	0.44%	
Regime 3 (Bull)	1.12%	1.49%	0.14%	
Regime 4 (Recovery)	3.61%	6.53%	1.03%	
Correlations/Volatilities Matrices				
Regime 1 (Crash) – Large Caps	6.09%			
Small Caps	76.04%	9.41%		
LT Bonds	-3.50%	-3.78%	2.59%	
Regime 2 (Slow Growth) – Large Caps	2.91%			
Small Caps	73.34%	4.11%		
LT Bonds	12.25%	-4.15%	1.87%	
Regime 3 (Bull) – Large Caps	3.25%			
Small Caps	73.19%	3.69%		
LT Bonds	2.79%	-1.49%	0.83%	
Regime 4 (Recovery) - Large Caps	3.44%			
Small Caps	27.10%	3.64%		
LT Bonds	31.40%	-22.52%	3.30%	
Transition Probabilities	Regime 1	Regime 2	Regime 3	Regime 4
Regime 1 (Crash)	77.48%	0.00%	0.00%	22.52%
Regime 2 (Slow Growth)	5.91%	91.44%	0.00%	2.65%
Regime 3 (Bull)	1.43%	0.00%	98.57%	0.00%
Regime 4 (Recovery)	8.40%	27.78%	3.68%	60.14%

Table 2: Probability Volatility Matrix and Ergodic Probabilities. We show below the estimated volatility matrix for the probability processes when there are 4 regimes, three risky assets and no predictability. The estimation procedure is explained in Section 5.1. Data are from January 1953 to July 2018. The table also reports the ergodic (unconditional) probabilities and the average duration of each regime.

	Z_1	Z_2	Z_3	Steady-State (Ergodic) Probabilities	Average Duration
Regime 1	-5.36%	-1.86%	0.06%	16.93%	4 months
Regime 2	2.48%	-1.23%	-1.21%	39.56%	12 months
Regime 3	0.84%	0.92%	0.10%	31.32%	70 months
Regime 4	2.04%	2.17%	1.04%	12.20%	3 months

Table 3: Dynamic Strategies and Approximate Solution Accuracy. We present the dynamic portfolio strategies for the four-state model without predictability, and the approximate solution accuracy measured by the total wealth equivalent loss. The optimal shares are obtained numerically by simulation, while the approximate ones are computed with the approximate analytical formulas presented in the text. Results are reported for different horizons, and portfolios are unconstrained, *i.e.*, short-selling is allowed. We consider that the next period regime can either be known (the investor sets the filtered probability $\mathsf{P}(\mathsf{Y}_{t+1} = i|\mathfrak{F}_t)$ to 1) or unknown (the investor uses the steady-state probabilities of the regimes).

Horizon (months)	Dynamic Portfolio Strategies								Total Wealth Equivalent Loss
	Optimal Shares				Approximate Shares				
	Large Caps	Small Caps	LT Bonds	Risk Free	Large Caps	Small Caps	LT Bonds	Risk Free	
Regime 1 (Crash)									
1	-78.0%	-9.0%	101.0%	86.0%	-76.5%	-7.5%	99.7%	84.2%	0.00%
6	-76.0%	-8.0%	96.0%	88.0%	-74.0%	-7.1%	98.4%	82.7%	0.00%
12	-76.0%	-8.0%	96.0%	88.0%	-73.5%	-6.8%	98.3%	82.0%	0.00%
60	-76.0%	-8.0%	97.0%	87.0%	-73.7%	-6.3%	98.5%	81.5%	0.01%
120	-76.0%	-8.0%	96.0%	88.0%	-73.7%	-6.3%	98.5%	81.5%	0.01%
Regime 2 (Slow Growth)									
1	252.0%	-115.0%	-8.0%	-29.0%	252.5%	-114.7%	-7.6%	-30.1%	0.00%
6	252.0%	-116.0%	-7.0%	-29.0%	251.6%	-115.5%	-7.9%	-28.2%	0.00%
12	252.0%	-115.0%	-7.0%	-30.0%	251.9%	-115.0%	-7.9%	-29.0%	0.00%
60	252.0%	-113.0%	-7.0%	-32.0%	251.5%	-113.7%	-7.6%	-30.3%	0.00%
120	252.0%	-113.0%	-6.0%	-33.0%	251.5%	-113.5%	-7.5%	-30.4%	0.00%
Regime 3 (Bull)									
1	24.0%	147.0%	-598.0%	527.0%	23.6%	145.9%	-599.8%	530.3%	0.00%
6	24.0%	143.0%	-603.0%	536.0%	22.9%	140.9%	-604.1%	540.4%	0.00%
12	24.0%	141.0%	-604.0%	539.0%	23.6%	139.6%	-605.1%	542.0%	0.00%
60	26.0%	137.0%	-607.0%	544.0%	24.9%	135.9%	-607.3%	546.5%	0.00%
120	26.0%	136.0%	-609.0%	547.0%	25.1%	135.4%	-607.6%	547.2%	0.00%
Regime 4 (Recovery)									
1	182.0%	928.0%	288.0%	-1298.0%	180.2%	930.1%	291.0%	-1301.3%	0.00%
6	182.0%	926.0%	288.0%	-1296.0%	179.7%	923.8%	287.7%	-1291.2%	0.00%
12	182.0%	926.0%	286.0%	-1294.0%	180.6%	924.1%	287.6%	-1292.3%	0.01%
60	182.0%	923.0%	284.0%	-1289.0%	181.0%	925.0%	287.7%	-1293.7%	0.01%
120	182.0%	923.0%	284.0%	-1289.0%	181.1%	925.1%	287.8%	-1293.9%	0.02%
Steady-State (Ergodic) Probabilities									
1	34.0%	49.0%	63.0%	-46.0%	35.3%	50.3%	64.0%	-49.6%	0.00%
6	35.0%	47.0%	62.0%	-44.0%	35.9%	48.2%	62.3%	-46.4%	0.00%
12	36.0%	47.0%	62.0%	-45.0%	36.4%	48.2%	62.1%	-46.7%	0.00%
60	36.0%	47.0%	62.0%	-45.0%	36.6%	48.0%	62.0%	-46.6%	0.01%
120	36.0%	47.0%	62.0%	-45.0%	36.6%	48.0%	61.9%	-46.6%	0.01%

Table 4: Expected Loss of a Myopic Investor. We report the total expected wealth equivalent loss (evaluated at the beginning of the investment period) of an investor following a myopic dynamic strategy as a percentage of the optimal dynamic strategy with hedging demands. We consider three different horizons of one, ten and forty years (given in months) and five cases: the four cases where the initial regime is known and a fifth case where the initial regime is unknown and the four regimes are weighted with the ergodic probabilities. Panel A reports the results for the no predictability cases. In Panel B, with predictability, the predictor value is set at its sample average and at this value plus or minus two standard deviations. The relative risk aversion of the investor is set at $\gamma = 5$.

Horizon	Regime 1	Regime 2	Regime 3	Regime 4	Ergodic
Panel A: Four-State Model Without Predictability					
T = 12	0.01%	0.00%	0.03%	0.03%	0.01%
T = 120	0.07%	0.06%	0.29%	0.08%	0.10%
T = 480	0.41%	0.35%	0.59%	0.46%	0.40%
Panel B: Four-State Model With Predictability					
$\mathbf{p_t = \overline{p} + 2\delta_{p_t}}$					
T = 12	0.02%	0.01%	0.02%	0.04%	0.00%
T = 120	1.23%	1.74%	23.86%	10.72%	5.72%
T = 480	88.10%	69.60%	98.92%	76.97%	90.85%
$\mathbf{p_t = \overline{p}}$					
T = 12	0.00%	0.02%	0.01%	0.04%	0.00%
T = 120	1.34%	1.72%	23.81%	10.81%	5.68%
T = 480	88.21%	69.73%	98.94%	77.18%	90.92%
$\mathbf{p_t = \overline{p} - 2\delta_{p_t}}$					
T = 12	0.00%	0.04%	0.01%	0.04%	0.01%
T = 120	1.32%	1.59%	23.85%	10.77%	5.74%
T = 480	88.31%	69.87%	98.95%	77.37%	91.00%

Table 5: Out-of-sample Performance for the Four-Asset Model. We compare the certainty-equivalent of the out-of-sample returns obtained with three portfolio models: the four-regime model, the single-regime model and the equally-weighted model. The horizon is set to 10 years and we rebalance the portfolios monthly. For the four-regime and single-regime models we computed the optimal portfolios with the risk aversion parameter $\gamma = 5$. To give a broader perspective of the performance over time, we start the out-of-sample exercise at different times, namely January 1960, January 1980 and January 2000. Samples end in July 2018. Panel A reports the first four unconditional moments of the returns for each strategy as well as the autocorrelation of order 1 and the turnover (monthly average of the traded shares divided by the portfolio size. Panel B reports the differences in the annualized certainty-equivalent return (ΔCER) between the four-regime and the two competing models with 95% bias-corrected and accelerated bootstrap confidence intervals.

Panel A: Descriptive Statistics of Out-of-Sample Returns

Portfolios	Mean (%)	Standard Deviation (%)	Skewness	Kurtosis	Auto-correlation (%)	Turnover (%)
Since 1960						
Four-Regime	1.56	7.94	-0.38	16.02	11.30	27.20
Single-Regime	0.83	3.09	-0.30	1.30	11.21	1.37
Equally-Weighted	0.73	2.48	-0.37	1.53	12.82	1.24
Since 1980						
Four-Regime	1.83	7.83	0.27	17.83	12.13	27.27
Single-Regime	0.92	3.01	-0.39	1.54	11.10	1.41
Equally-Weighted	0.78	2.41	-0.55	1.89	12.94	1.26
Since 2000						
Four-Regime	1.39	6.31	2.82	25.89	-1.59	27.27
Single-Regime	0.63	2.71	-0.49	1.89	5.98	1.41
Equally-Weighted	0.51	2.29	-0.43	1.47	8.49	1.26

Panel B: Δ CER	Since 1960	Since 1980	Since 2000
Four-Regime *vs.* Single-Regime	-0.069 [-0.216 – -0.008]	-0.096 [-0.171 – -0.091]	0.069 [0.056 – 0.145]
Four-Regime *vs.* Equally-Weighted	–0.064 [-0.177 – -0.001]	-0.099 [-0.152 – -0.090]	0.073 [0.056 – 0.182]

Table 6: Parameter Estimates of the Four-Regime Model with Predictability. We report below the parameters for the four state model with predictability. The matrices $\mathbb{A}$ and $\mathbb{B}$ are defined in equation (5) in the main text. The volatilities are reported on the diagonals of the correlations/volatilities matrices between the assets. Just below, we include the transition and ergodic probabilities of the four-regime model in the presence of a predictor. The ergodic mean values in each regime and the volatilities of the predictor (see equation (3) in the main text) are reported at the bottom of the table, together with the predictor's speed of mean reversion. The estimation period is from January 1953 to July 2018. The average Treasury yield over the full period was estimated at 0.36% a month.

Four-State Model with Predictability	Large Caps	Small Caps	LT Bonds
Matrix $\mathbb{A}$ Transposed - Regime 1 (Crash)	-2.98%	-4.24%	0.46%
Regime 2 (Slow Growth)	1.03%	0.69%	-0.26%
Regime 3 (Bull)	0.99%	1.32%	-0.15%
Regime 4 (Recovery)	3.87%	6.53%	0.64%
Vector $\mathbb{B}$ Transposed	3.59%	5.36%	7.74%
Mean Returns - Regime 1 (Crash)	-2.70%	-3.82%	1.08%
Regime 2 (Slow Growth)	1.11%	0.81%	0.44%
Regime 3 (Bull)	1.16%	1.57%	0.22%
Regime 4 (Recovery)	3.82%	6.45%	0.52%
Correlations/Volatilities Matrix			
Regime 1 (Crash) - Large Caps	4.97%		
Small Caps	80.12%	7.65%	
LT Bonds	-6.56%	-1.92%	2.61%
Regime 2 (Slow Growth) - Large Caps	2.91%		
Small Caps	72.95%	4.14%	
LT Bonds	16.39%	-3.18%	1.89%
Regime 3 (Bull) - Large Caps	3.21%		
Small Caps	72.18%	3.66%	
LT Bonds	2.58%	-1.99%	0.83%
Regime 4 (Recovery) - Large Caps	3.85%		
Small Caps	17.91%	5.25%	
LT Bonds	24.17%	-14.65%	3.13%

Probabilities	Regime 1	Regime 2	Regime 3	Regime 4
Transition Prob. - Regime 1 (Crash)	61.21%	0.26%	0.01%	38.52%
Regime 2 (Slow Growth)	7.28%	92.68%	0.01%	0.03%
Regime 3 (Bull)	1.39%	0.00%	98.61%	0.00%
Regime 4 (Recovery)	19.39%	19.65%	3.19%	57.77%
Steady State (Ergodic) Prob.	15.34%	38.23%	32.40%	14.02%

Predictor Parameters	Ergodic Mean ($\bar{\mathbb{P}}$)	Volatility Vector (σ_p)		Mean-Reversion Speed (κ_p)
Regime 1 (Crash)	0.0792	σ_{p1}	-0.0005	0.0163
Regime 2 (Slow Growth)	0.0223	σ_{p2}	0.0000	
Regime 3 (Bull)	0.0471	σ_{p3}	-0.0002	
Regime 4 (Recovery)	-0.0155	σ_{pp}	-0.0022	

E Figures

Figure 1: Dynamic Strategies for Different Horizons in the Model With Predictability. We plot the strategies for horizons from one month to ten years. We consider the case where the current regime is unknown and set the current probabilities at their ergodic values for each state. Each panel corresponds to a class of assets, respectively large stocks, small stocks, long-term bonds and Treasury bills. In each panel, three curves are drawn, one when the starting value of the predictor is set at its mean, the two other ones when it is set at its mean plus or minus two standard deviations. The investor's relative risk aversion γ is set at a value of 5.

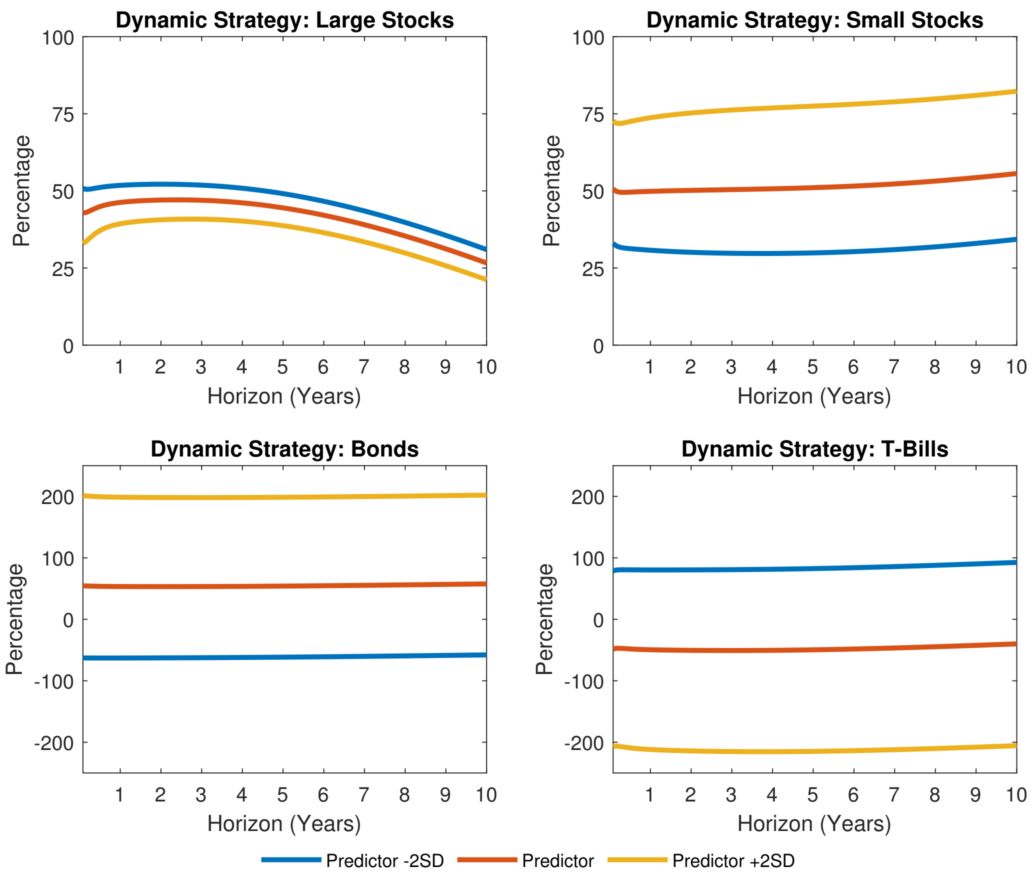

Figure 2: Consumption to Wealth Ratio. We plot below the consumption strategy for dynamic investors with $\gamma = 5$. In the upper plots, we draw six curves depending on the current regime for horizons from one month to 10 years ($\psi = 0.75$). In the lower plots, we show the consumption strategy followed by a 12-month horizon investor using real-time filtered probabilities (and predictor values) for the period from July, 1988 to July, 2018. The left plots are for the problem with no predictability while the ones in the right, for the model with predictability.