Optimal Goals-Based Investment Strategies For Switching Between Bull and Bear Markets

Revision 1, August 2021

Authors:

1. Sanjiv R. Das (Professor of Finance and Data Science at Santa Clara University), 500 El Camino Real, Santa Clara, CA 95053. srdas@scu.edu
2. Daniel Ostrov (Professor of Mathematics at Santa Clara University), 500 El Camino Real, Santa Clara, CA 95053. dostrov@scu.edu
3. Aviva Casanova (Manager, Client Insights & Analytics at Franklin Templeton), One Franklin Parkway San Mateo, CA 94403. aviva.casanova@franklintempleton.com
4. Anand Radhakrishnan (Director, Strategy and Analytics at Franklin Templeton), One Franklin Parkway San Mateo, CA 94403. andy.radhakrishnan@franklintempleton.com
5. Deep Srivastav (SVP, Head of Digital Solutions at Franklin Templeton), One Franklin Parkway San Mateo, CA 94403. deepratna.srivastav@franklintempleton.com

The first author is the corresponding author.

Optimal Goals-Based Investment Strategies For Switching Between Bull and Bear Markets

September 4, 2021

Abstract

We solve a dynamic long-horizon goals-based wealth management problem, given different investment regimes. In a world with a good regime (bull market) and a bad regime (bear market), an investor who is cognizant that regime switching occurs has the potential to do better than an investor who assumes only one regime. However, models with more than one regime incur the additional risk of regime uncertainty. Investors must be able to predict which regime is governing the market with reasonable levels of confidence, or they can be worse off than investors who assume just one regime. Using data from recent history, we develop a framework that determines how accurate regime prediction needs to be in order to achieve gains from a regime-cognizant goals-based investing approach.

Highlights

1. An optimal dynamic program for goal-based investing when switching between efficient frontiers in bull and bear markets.
2. Dealing with regime risk via endogenous Bayesian updating and using exogenous external information.
3. The model ascertains the accuracy level of regime prediction needed for a regime-switching model to perform better than a single-regime model.

Keywords: Regimes, goals-based investing, dynamic optimization, uncertainty
JEL codes: G02, G11

1 Introduction

Investors tend to buy stock in bull market regimes and sell stock in bear market regimes. This well known behavior is illustrated in Exhibit 1, which shows the strong positive correlation between stock market performance (as represented by the S&P 500 index) and monthly cash flows from investors in U.S. stock mutual funds. In this paper, we analyze fund choice for goals-based investors who are managing wealth over long horizons in a regime-changing world. Not surprisingly, the optimal investment strategy is quite different from this typical investor behavior.

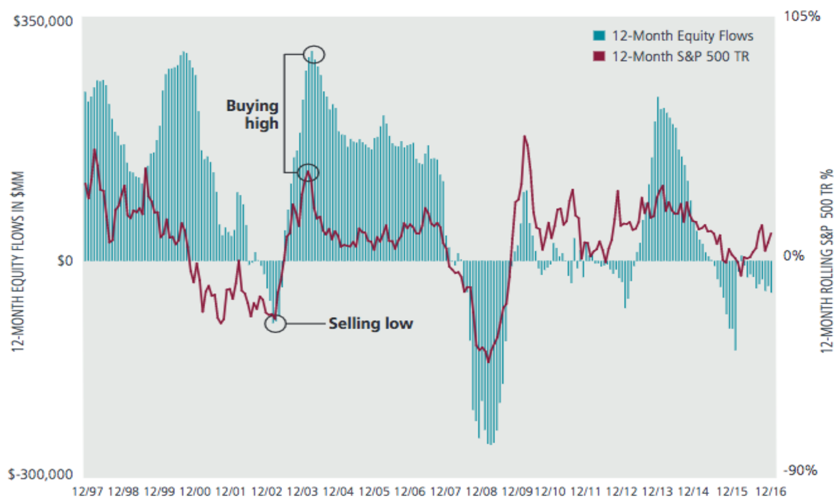

Exhibit 1: The strong tendency of investors to move money into stocks in bull market regimes and out of stocks in bear market regimes. Source: http://kattanferrettifinancial.com/buying-high-selling-low/.

We consider investors with horizons over which regimes may switch between bull and bear markets. We assume that investors aim to reach a target level of wealth at a future date, an approach now well-known as goals-based investing, see Chhabra (2005), Brunel (2015). This paper assesses how goals-based investors might dynamically manage a portfolio over broad asset classes in the presence of regime switching.

Observing the investor behavior in Exhibit 1, where investors move in and out of stock (switching, presumably, to bonds primarily) as they react to regime switches, raises a number of natural questions that we look to answer in this paper: (i) The investment analyses of regime-switching are often limited to just a stock and a bond position, but what is the optimal investment philosophy if we have multiple assets to consider? (ii) In contrast to extant investor behavior, but in line with most professional financial advice, is the optimal investment strategy to become more, not less, aggressive in bear markets and more conservative in bull markets, and, if so, by how much? (iii) How do we take into account the fact that investments behave differently in a bull market regime versus a bear market regime? For example, during a bear market, the volatility of stocks generally increases, and the returns for different stocks become more correlated. (iv)

How do we best define different regimes for stock market behavior? (v) Should we consider more than just two regimes? Are we better off using a single regime instead of two regimes? (vi) How do we handle the uncertainty about knowing which regime the investor is in? (vii) How sensitive are the outcomes for goals-based investing to the data used for regime estimation?

The answers given in this paper to these questions come from combining ideas from a number of different fields of study. Econometrics is used to determine the best way to split market behavior into two (or more) different regimes. Dynamic programming determines the optimal investment strategy within this multiple regime model. This is a large extension to standard single-regime dynamic optimization models. Goals-based wealth management (GBWM) with dynamic programming optimizes the probability of reaching the investor's goal. Bayesian updating determines how the investor, moving forwards in time, will decide which regime to assume while making optimal fund choices. We also determine how external information can be balanced with Bayesian updating for better regime prediction.

Brief answers to the questions above are as follows: (i) Our model easily accommodates multiple assets and always chooses portfolios (funds) formed from these assets that reside on the regime-specific efficient portfolio frontier. Therefore, the approach is one where we dynamically slide up and down an efficient frontier or switch efficient frontiers based on the state (wealth, time, regime) of the portfolio. (ii) Within a specific regime, we find that it is better to take some more risk when the portfolio is underperforming, though the amount of extra risk is different in the two regimes. The exact amount is computed in the algorithm, making it more precise than rules of thumb used in practice. (iii) In order to take account of the differences between regimes, we generate efficient frontiers for both the good and bad regimes. These are based on all available assets chosen for the model. The optimal fund along these frontiers is then chosen using dynamic programming to determine the optimal strategy over time. (iv) We determine the two regimes by applying the regime-switching econometric methods in Hamilton (1989) to the returns of a balanced index fund. This allows both equity and bond markets to be represented in the time-series used to determine when we are in each of the two regimes. We find that about one-fourth of the time is spent in the bad regime. (v) We note that adding a third regime to the mix does not improve a variety of statistical values that measure the trade off between the number of regimes and goodness-of-fit with the time series of the underlying index. Therefore, we choose a regime-switching model with exactly two regimes. The answer to the question of whether it is better to use two regimes or one regime is more complicated and discussed in the next paragraph. (vi) Finally, as expected, the empirical outcomes are highly sensitive to the model's inputs, making it key that investors pay attention to the period of time over which they estimate the parameters of the two regimes. We present some examples of this sensitivity using data from subsamples, showing that the model's performance depends significantly on the estimated inputs. This difference is generally quantitative, but it can be qualitative as well, for example, changing whether it is better or worse to take more risk after a transition from the good regime to the bad regime. This highlights that the forecasting model used to determine the algorithm's inputs is an important modeling consideration.

For comparison with single regime dynamic optimization models, we recognize that investors experience three sources of uncertainty in markets that can move between regimes. First, the traditional source of risk is the random value drawn from the stochastic model for the returns of a chosen fund if we know the governing market regime; more specifically, this is usually the effect from a random variable that is multiplied by the volatility assigned to the fund. This risk also exists in single regime models. Second, in multiple regime models, a source of risk comes from not knowing precisely which regime the investor will experience in each period. We will see that if an investor has weak knowledge of which regime is extant, the investor may in fact be better off assuming a one-regime model. Our framework allows us to use information about the extant regime that is exogenous to the data used in the fund optimization. This will allow us to ascertain the level of exogenous regime knowledge needed to make the two-regime model superior to a one-regime model for a goals-based investor. A third source of risk comes from return distribution parameter instability in the model over time (nonstationarity). Estimates used to determine future period funds and strategies are fraught with forecast error, and this risk may be higher in multiple regime models versus single regime ones.

The rest of the paper proceeds as follows: In Section 2 we review the investment literature that considers changes to the underlying economic regime. Section 3 presents our methodology, including the dynamic programming-based approach we apply in the paper. Section 4 describes and analyzes results. Section 5 contains concluding comments.

2 Risk Regimes

Before we discuss this paper's approach to the goals-based wealth management (GBWM) problem with regimes, we undertake a brief discussion of the relevant literature. Early literature on asset allocation, starting from Markowitz (1952), evolved into a literature on asset pricing with changing states of the world, as in, for example, Merton (1973). A special case of the Merton paper encompasses changes in state due to regime switching. An extensive econometric literature centering on regime-switching models started with the seminal paper by Hamilton (1989). In Zhou and Yin (2003), portfolio optimization with regime switching is explored with the goal of minimizing the expected squared difference between a given terminal wealth and the actual terminal wealth.

More recently, the literature on asset allocation with regime switching has considered optimizing lifetime expected utility, summarized in Ang and Timmermann (2012). History has shown that financial markets swing suddenly from good economic regimes to bad ones and vice versa, and they display very different behavior in each regime. (Theoretically, this may be the outcome of behavioral dynamics, Branch and Evans (2010).) Bad economic regimes typically have lower returns and higher volatility. They also have more highly correlated returns, resulting in a loss of portfolio diversification when it is most needed; see Das and Uppal (2004), Singhal and Biswal (2019). Therefore, having a model that captures these regime changes can be quite helpful in managing wealth in long-horizon portfolios, especially since regime changes are quite common. In the main illustration in this paper, for example, we see that over the past two decades, our model shows the average time in a regime to be between 5 and 17 months. Ang and Bekaert (2002) show that regime effects exist in international asset markets, and the impact of ignoring them is significant when a risk-free asset is available. Pagan and Sossounov (2003) and Guidolin and Timmermann (2008) also show that there are distinct bull and bear regimes in international stock markets. Ammann and Verhofen (2006) find that value stocks do better in the high-variance regime, whereas, in the low-variance regime, momentum stocks do better. They find that regime forecasting ability is weak, as we do. Hu (2019) suggests that it is important to account for regimes in portfolio decisions, but accurate determination of the structure and number of regimes is of equal importance.

In practice, the wealth management industry has generally handled regime changes as they occur, using static optimization in a purely myopic manner. Theoretically, the literature has mostly shown that dynamic optimization with changing states is considerably better, resulting in gains of 200 basis points and more over static myopic strategies; see Tu (2010). Ang and Bekaert (2004) show that, for global all-equity portfolios, a regime-switching strategy dominates myopic strategies in an out-of-sample test; see also Bernhart et al. (2011), Bulla et al. (2011), Grobys (2012), Jiang et al. (2015), Dapena et al. (2019), and Lewin and Campani (2020) for similar results. Guidolin and Timmermann (2007) consider four regimes (crash, slow growth, bull, and recovery) that exist across equity and bond markets, and they show that optimal asset allocations differ drastically across these regimes. In poor states, they conclude that investors should take more risk when their horizons are long, a finding we corroborate as a byproduct of our analysis. Regime-switching has also been analyzed in the presence of inflation and short-selling (Bellalah et al., 2020).

We complement this literature by considering fund optimization with switching regimes in a goals-based context, which has not yet been explored in wealth management research. The traditional notion of portfolio optimization is a trade-off between risk and return in the portfolio's assets (Zhou and Yin, 2003), where risk is measured by the standard deviation of the portfolio's returns (that is, the portfolio volatility), or in a utility framework with downside constraints (Vo and Maurer, 2013). However, in a goals-based context, portfolio risk is the probability that an investor does not meet their goals, and this is the notion of risk we wish to minimize. This goals-based notion of risk is different from the volatility notion of risk, since a young investor saving for retirement would have a high risk of not achieving their retirement goals were they to move all of their retirement money into a low volatility risk asset like cash. We note that the goals-based view of risk still leads us to remain on the efficient frontier, where we must be for the trade-off between traditional risk (volatility) and expected return to be optimized. Evidence presented in Das et al. (2018) shows that investors are more comfortable with the notion of the probability of meeting goals than with mathematical concepts such as Sharpe ratios, correlation measures of diversification, etc. Goals-based wealth management was introduced by Nevins (2004) and Chhabra (2005), and an exhaustive description of the paradigm is undertaken in Brunel (2015). Regime switching most certainly impacts the investor's probability of reaching goals, and we work towards developing the optimal strategy for an investor to pursue in the presence of these regime shifts. Inevitably, shifting into a bad economic regime might trigger adverse responses if goals are not kept in mind, and it will also dictate a different strategy from the case where the regime switch did not occur. We will compare how much difference regime switching makes to a goals-based strategy. Such an analysis has not so far been considered in the regime-switching literature nor in goals-based wealth management models. Interestingly, unlike the optimistic literature on asset management with regimes, this paper finds that regime estimation error, if not minimized, offers a poor prognosis for wealth management outcomes over long horizons in the presence of regime switching. Assuming a single regime model is generally best in such circumstances, in the spirit of simpler portfolio prescriptions such as those in DeMiguel et al. (2009). However, we will be able to quantify how accurate regime estimation methods need to be in order for a two-regime model to again yield superior advice to the advice determined with a single regime.

3 Regime Methodology

In this section we describe our approach to dynamically optimizing investments in a two regime world so as to maximize the probability of attaining a goal wealth of G dollars or more at the portfolio's horizon time $t = T$. In Subsection 3.1, we use the historical results from a single index to determine which of the two regimes was most likely controlling the economy at each time in that history. From this we are able to determine the different efficient frontiers for each of the two regimes using as many assets as desired, although we will use three assets in our examples, representing the major asset classes used by target date funds. Subsection 3.2 summarizes the technical details contained in the Appendix A regarding the regime dynamics, our dynamic programming approach, and the incorporation of internal and external information for determining the governing regime.

3.1 Determining the regime transitions and regime-specific efficient frontiers

We estimate the two economic regimes using monthly returns from a single index for bond and stock markets over the period of ~ 17 years from February 2004 through March 2021. We used the Vanguard Balanced Index Fund (VBINX), which is 60% U.S. stocks and 40% U.S. bonds, for the index, and note that almost identical regimes were obtained when we used the Vanguard Global Equity Fund (VHGEX) instead.

Our econometric model for regime switching assumes a different constant mean and standard deviation for the index's returns in each of the two regimes. For the good regime, we will have μ_0 and σ_0 as the mean and standard deviation; the corresponding parameters for the bad economic regime will be μ_1 and σ_1. We also assume constant monthly transition probabilities between the two regimes; specifically, the constant p_{01} denotes the probability of transitioning from Regime 0 to Regime 1, and the constant p_{10} is the reverse transition probability. Of course, we stay in Regime 0 with probability $p_{00} = 1 - p_{01}$, and we stay in Regime 1 with probability $p_{11} = 1 - p_{10}$. This notation corresponds to the elements of the regime transition probability matrix

$$\mathbf{p} = \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix}$$

that defines the random walk regime transition dynamics from time t to time $t+1$. We note that Pagan and Sossounov (2003) find that a pure random walk model provides as good an explanation of bull and bear markets as more complex statistical models.

Within each regime, the returns, Ret_t, for the Vanguard Balanced Index Fund at each time t are modeled by

$$Ret_t = \mu_i + \epsilon_{i,t}, \quad \epsilon_{i,t} \sim N(0, \sigma_i^2),$$

where $i = 0$ or 1, depending on which economic regime we have at time t. Given the history of returns, Ret_t, we look to estimate the parameters for this model, specifically, the mean return in each regime, μ_0, μ_1, the standard deviation of return in each regime, σ_0, σ_1, and the transition probabilities, p_{00}, p_{10}. (The remaining two transition probabilities follow from $p_{01} = 1 - p_{00}$ and $p_{11} = 1 - p_{10}$.)

Estimation of these parameters using maximum likelihood estimation is standard (see, for example, Kim and Nelson (1999), who use Gibbs sampling, and Kupelian (2020), who uses change point analysis). We implement this using the `statsmodels` package for the Python programming language and the `MarkovRegression` function. The inputs to this package are all 206 monthly returns, Ret_t, from the ~ 17 year period of collected return data for the Vanguard Balanced Index Fund. The resulting estimated parameters are shown in Exhibit 2.

The probability of being in each regime over time is also estimated. Exhibit 3 compares the estimated probability of being in Regime 0 over time to the returns, Ret_t, over time.

We annualize the parameters in Exhibit 2 by multiplying the mean by 12 and the standard deviation by $\sqrt{12}$. In the good regime, the annualized mean return is 10.70% with a standard deviation of 5.50%. In the poor regime the annualized mean and standard deviation are 0.66% and 18.52% per annum. Ignoring regimes, the annualized mean and standard deviation of returns across the entire ~ 17 year time series are 8.19% and 10.53% per annum, respectively. The market remains in the good regime for an average of $\frac{1}{p_{01}} = \frac{1}{0.0596} = 16.77$ months and in the bad regime for an average of $\frac{1}{p_{10}} = \frac{1}{0.1837} = 5.44$ months. This implies that we see calm markets (the good regime) 3/4 of the time and low return, volatile markets (the bad regime) 1/4 of the time.

Exhibit 2: Estimated parameters for the regime-switching model. Two regimes are estimated, Regime 0 (the good regime, e.g., a bull market) and Regime 1 (the bad regime, e.g., a bear market). The units of time in the results are in months, reflecting the fact that returns are monthly. The regime transition probability matrix $\mathbf{p}$ is also shown. The data used for detecting the regimes are the monthly returns from the Vanguard Balanced Index fund (VBINX), from February 2004 through March 2021.

Regime	Parameter	Value	Std Error
Regime 0 (Good)	μ_0	0.8920%	0.141%
	σ_0	1.5880%	
Regime 1 (Bad)	μ_1	0.0552%	0.776%
	σ_1	5.3476%	
Transition	p_{00}	0.9404	0.025
Probability	p_{10}	0.1837	0.075

$$\mathbf{p} = \begin{bmatrix} 0.9404 & 0.0596 \\ 0.1837 & 0.8163 \end{bmatrix}$$

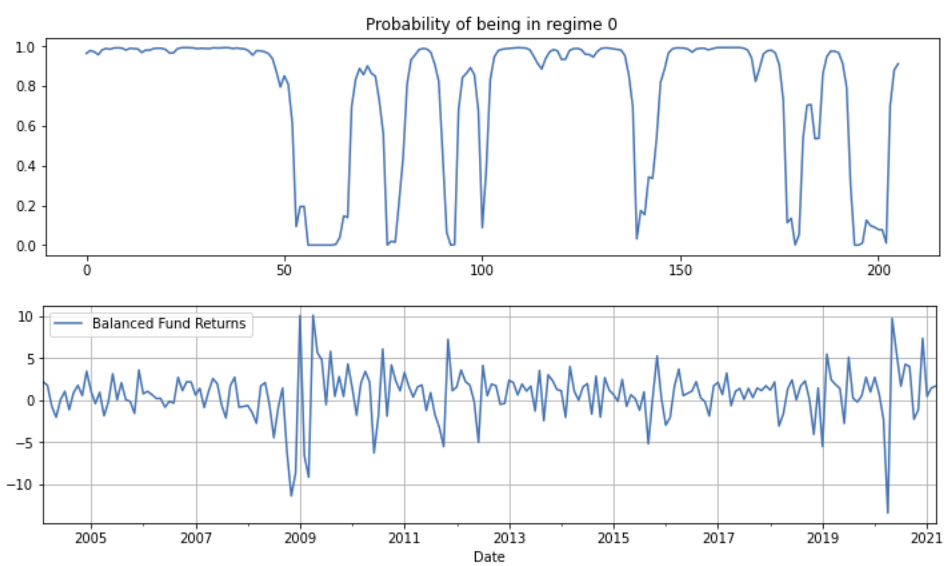

Exhibit 3: Top panel: The probability of being in Regime 0 (the good regime) over time. Bottom panel: The returns, Ret_t, for the market over time. Over the displayed ~ 17 year time frame (02/2004 to 03/2021), the monthly mean excess return is 0.81%, with a standard deviation of 4.32%. Comparing these two panels, we see that the good regime seems to coincide with periods of low volatility.

To determine the properties of portfolio assets in the good regime and the bad regime, we assume that the good regime corresponds to all times when we are more likely in the good regime than the bad regime; that is, all the times in the top panel of Exhibit 3 when the curve is greater than 0.5. Otherwise we assume we are in the bad regime. Given this assumption, we constructed the efficient frontier in each regime for the following three assets from return data over the same period used previously: (1) the Vanguard Large-Cap Index Fund (VLACX), (2) the Vanguard Small-Cap Index Fund (VSMAX), and (3) the Vanguard Total Bond Market Index Fund (VBMFX). The two efficient frontiers, along with the mean returns and covariances of returns within each regime for the three assets, are shown in Exhibit 4.

Exhibit 4 also shows the efficient frontier, the mean returns, and covariances of returns if we assume that there is only a single regime. This frontier lies between the frontiers for the two regimes, which may seem intuitive, but this is not generally the case, as we will see, for example, in the cases considered in Appendix B. The efficient frontier in the good regime dominates the frontiers in both the poor regime and the single regime, except for a very small range in the lowest risk region. There, the poor regime offers better bond returns than the good regime, with only marginally higher volatility (4.8%) over the single-regime (3.9%) and good regime cases (3.6%). This is understandable, since bonds are likely to perform better in poor regimes.

3.2 Portfolio algorithm for two regimes

A brief overview of the algorithm for optimizing goal-based portfolios in a two-regime world is described here, relegating the technical details to Appendix A. Without loss of generality, we assume that portfolio wealth is driven by geometric Brownian motion with a mean and volatility determined by the fund chosen along the governing regime's efficient frontier. Given the portfolio's time horizon T, we use dynamic programming to determine which fund on the governing regime's efficient frontier is optimal each month for an investor looking to maximize the probability of reaching their goal wealth, G. This optimal fund choice will depend on the remaining time to the horizon, $(T-t)$; the current portfolio wealth, W; and the current governing economic regime, r_0 (the good regime) or r_1 (the bad regime). This optimal fund solution is obtained using backward recursion, which assumes perfect knowledge of the regime. Because we assume perfect knowledge of the regime, this backwards recursion pass gives the highest goal probability possible.

In order to accommodate regime uncertainty, the algorithm next implements a forward pass by generating several simulated paths of regime transitions and asset returns. The fund chosen at each time and wealth is taken from the backwards recursion pass's result for the optimal fund choice, but now using a guess for the current governing economic regime. This fund choice will, of course, be suboptimal if the guess for the regime is incorrect. In the simulated forward pass, the guess for the regime in each period is determined from a weighted combination of Bayesian updating along the sample path of returns (internal information) and injecting external information of varying accuracy about the true regime. Of course, when the accuracy of external information

Exhibit 4: The efficient frontiers in both regimes (Regime 0, the good regime, and Regime 1, the bad regime), and the efficient frontier assuming a single regime. We used daily data and then annualized the return values. We use (1) the Vanguard Large-Cap Index Fund (VLACX), (2) the Vanguard Small-Cap Index Fund (VSMAX), and (3) the Vanguard Total Bond Market Index Fund (VBMFX). The calculated parameters used to determine these three efficient frontiers are contained in the table below the graph.

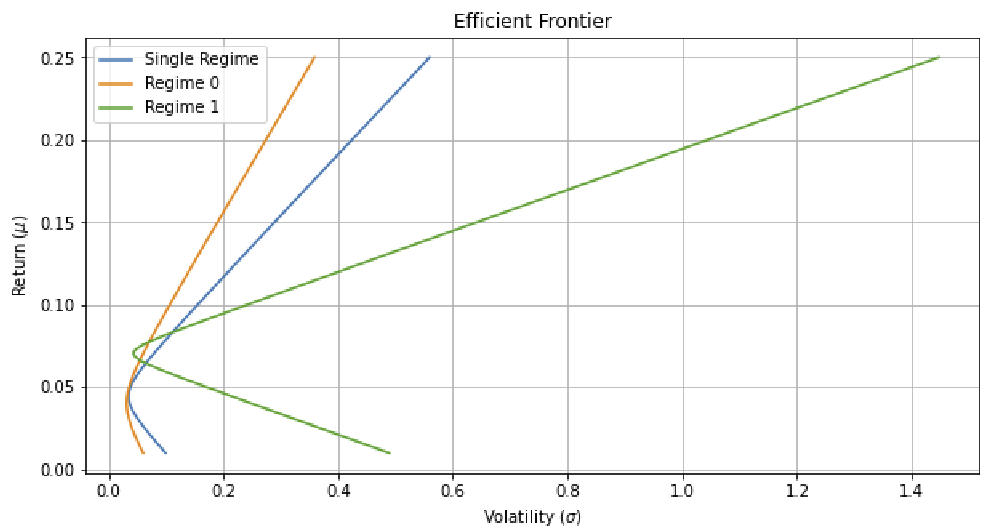

Single Regime	VLACX	VSMAX	VBMFX
Mean return (annual)	0.1137	0.1274	0.0378
Covariances (annual)			
VLACX	0.037595	0.042393	-0.002236
VSMAX	0.042393	0.054331	-0.002539
VBMFX	-0.002236	-0.002539	0.001538
Regime 0	VLACX	VSMAX	VBMFX
Mean return (annual)	0.1144	0.1298	0.0296
Covariances (annual)			
VLACX	0.017140	0.019660	-0.001337
VSMAX	0.019660	0.026655	-0.001571
VBMFX	-0.001337	-0.001571	0.001323
Regime 1	VLACX	VSMAX	VBMFX
Mean return (annual)	0.1114	0.1183	0.0678
Covariances (annual)			
VLACX	0.112564	0.125710	-0.005533
VSMAX	0.125710	0.155770	-0.006087
VBMFX	-0.005533	-0.006087	0.002322

about the regime is perfect and none of the internal information is used, the goal probability from the forward pass equals that from the backward recursion based dynamic programming solution.

The technical details of the approach just outlined are given in Appendix A. More specifically, in Subsection A.1 of the Appendix, we consider a set of possible funds along each efficient frontier and determine the probabilities of moving from each possible wealth value at time t to a possible wealth value at time $t+1$, given the regime and a possible fund in that regime. This is used in Subsection A.2 of the Appendix to formulate the dynamic programming model and its associated Bellman equation, which determines the optimal fund to hold at any known time, wealth value, and regime. Subsections A.3 and A.4 of the Appendix address the fact that the investor always knows the current time and their current wealth, but they can never be certain of their current regime. In Subsection A.3, we develop the Bayesian model for guessing the regime from the internal information. In Subsection A.4, we consider how to incorporate additional external information into our guess for the regime.

4 Analysis

In this section, we look to assess the performance of the dynamic programming-based approach to fund optimization with regime switching shown in Section 3. Past literature has assessed the performance of models that switch regimes and change funds using repeated static rebalancing. That is, for the time period starting at any given time t, the investor fits a regime-switching model to the data within an s-period historical window $(t-s,t)$ and then chooses a fund for the period that starts at time t based on a predicted regime and parameters estimated from the historical window. The entire exercise is then repeated for the period starting at time $t+1$, so re-estimation occurs every period, as undertaken, for example, in Ang and Bekaert (2004). Other papers (for example, Guidolin and Timmermann (2008)) assess the effect of higher-order moments and the presence of regimes on fund choices, which support observed home bias tilts in investor portfolios. This constant updating for parameters and the regime model every period in repeated static rebalancing is inconsistent with dynamic programming models for fund choice under regime switching, where future changes in regimes are assumed to come from a regime-switching model with constant parameters, and the dynamic program has already determined optimal fund choice.

This leads to a different experimental approach in this section that is based on how the theoretical analyses in Section 3 apply to the data. As previously discussed, we assume stationarity of the regime-switching system so that dynamically optimal fund decisions can be implemented. Our analyses in this section account for uncertainty in returns and regimes, with a view to comparing whether investors should continue using a single regime model for which, by its nature, there can be no regime uncertainty, versus a two-regime model for which, by its nature, there is regime uncertainty. In order to make this comparison, we do not assume recalibration each period so as to make a fair comparison of the single-regime and two-regime worlds. That is, it is assumed that the investor makes a forecast of the mean vector and covariance matrix of returns for each regime based on any forecast technology they choose. At the beginning of this section, we assume that this parameter forecast is accurate, so that we may focus on the effects of regime uncertainty in the determination of whether the investor should be using a multi-regime model or just one regime. At the end of the section, we will circle back to exploring questions about the accuracy of the forecasts for the parameters.

Whereas mean-variance optimization computes efficient frontiers in risk-return space, goal-based optimization maximizes goal probability using portfolios that are always mean-variance efficient, hence combining the best of the modern portfolio theory and behavioral portfolio theory of Shefrin and Statman (2000). This approach manages two types of risk: (i) standard risk (the volatility of portfolio returns) and (ii) behavioral risk (failure to meet goals), resulting in a more optimal approach in the spirit of Parker (2020) and Parker (2021).

In this section, we use parameters calibrated from the asset data in Subsection 3.1 to analyze the results from a variety of numerical experiments. First (Subsection 4.1), we will examine how different the optimal fund choices are in each regime, and, within each regime, we will quantify how the optimal fund choice changes from taking more risk to taking less risk as the portfolio changes from under-performing to over-performing to attain its goal wealth. Second (Subsection 4.2), we will show why the three-regime model is a worse choice than the two-regime model. Third (Subsection 4.3), we compare the two-regime model and the one-regime model in the context of regime certainty. That is, we explore our central question regarding which model is better. Fourth (Subsection 4.4), we explore our central question in the context of regime uncertainty, both by working with external information and considering a number of alternative models to Appendix A.3 for working with internal information. Fifth (Subsection 4.5), we will examine the effect of infusions and withdrawals into and out of the portfolio. Finally, sixth (Subsection 4.6), we consider uncertainty in the regime parameters themselves, as opposed to uncertainty in the regime in which we reside. Specifically, we will examine the robustness of the regimes' efficient frontiers to changing the historical time intervals used to determine them. This includes using subperiods from our return data history and slightly expanding the historical interval used for our return data.

4.1 Regime-based fund choice

We work with the example data from Subsection 3.1, where we used ~ 17 years of monthly data from February 2004 to March 2021 for a balanced index fund asset (VBINX) to determine the good and bad regimes in that period, and then used the returns of three index fund assets — (1) the Vanguard Large-Cap Index Fund (VLACX), (2) the Vanguard Small-Cap Index Fund (VSMAX), and (3) the Vanguard Total Bond Market Index Fund (VBMFX) — in the good regime and in the bad regime to determine the efficient frontier in each regime for these three index fund assets. We also determined the efficient frontier for these three index fund assets if there is only a single regime. Recall that all three of these efficient frontiers were shown in Exhibit 4.

For each of these three efficient frontiers, we select 15 potential funds, numbered 0 through 14, on the frontier. To avoid bias towards any regime, the set of means selected for the 15 funds is chosen to be the same, regardless of the regime, but of course, with different standard deviations for each regime. The specific index fund asset weights, the means, and the standard deviations for these 45 funds are shown in Exhibit 5.

We see that in almost all cases, the funds are long-only. Only in the bad regime do we have funds that short assets, specifically the funds numbered zero through four. However from Exhibit 4, we see that all five of these funds fall below the vertex of the hyperbola for the bad regime's efficient frontier, which is near $\mu = 0.07$. Since these funds fall below the vertex, they will always be dominated by funds on the efficient frontier above the vertex, so they are never used. In other words, all the funds used here are long-only.

We will consider the case of an investor with an initial wealth of \$100,000 whose goal is to grow this wealth to at least $G =$ \$225,000 after $T = 10$ years. Except for Subsection 4.5, we will assume that the portfolio experiences no infusions nor withdrawals during these 10 years. Using the dynamic programming method discussed in Appendix A.2, we determine Exhibit 6, which shows the optimal fund selection to attain this investor's goal given any time and the portfolio worth at that time in the case of a single regime (top panel) or, in the case of two regimes, if we're in Regime 0 (middle panel) or Regime 1 (bottom panel)[1] at the given time. Note that these optimal fund prescriptions assume regime certainty. In each of Exhibit 6's three panels, we see that more risk is taken when the wealth level is low and less risk is taken when the wealth level is high. This is intuitive — when the portfolio is not on track to reaching the goal, more risk is taken to reach for return, which improves the chance of meeting the goal. Likewise, when the portfolio is doing very well, risk is dialed back to help the portfolio stay on track towards attaining the investor's goal.

In other words, given that we stay in a specific regime, when the portfolio does worse, we are better off increasing risk, which is the opposite of the investor behavior we saw in Exhibit 1. But what happens when we switch regimes? As we can see in Exhibit 4, because the efficient frontier for the good regime generally dominates the efficient frontier for the bad regime, the risk-return tradeoffs in the good regime are generally better than those in the bad regime. This means that more risk is taken when the portfolio is underfunded in the bad regime than if it is underfunded in the good regime. This is not a feature of good versus bad regimes, but, instead it is a feature of which regime has a better risk-return tradeoff, indicated by which regime's efficient frontier lies above the other's. We will see examples where the bad regime lies above the good regime in Appendix B.

[1]Note in this panel for the bad regime that, as expected, we do not see the five funds that lie below the vertex of the bad regime's efficient frontier.

Exhibit 5: The 15 potential funds along each of three efficient frontiers shown in Exhibit 4. The three frontiers correspond to a single-regime model (top panel) and a two-regime model, namely, the good regime (middle panel) and the bad regime (bottom panel). The weights of the three component assets are given, along with the fund's mean and standard deviation. Only in the poor regime is shorting sometimes encountered for these asset classes, and the investor may choose to restrict it if necessary. These portfolios are dominated in any case. On purpose, the set of means for each regime's frontier have been chosen to be the same to avoid bias towards any regime.

Component Weights For Each Investment Fund Number

Single Regime

Fund Number	VLACX weight (Large-Cap)	VSMAX weight (Small-Cap)	VBMFX weight (Bonds)	Mean	Standard Deviation
0	0.150	0.010	0.839	0.050	0.038
1	0.183	0.035	0.782	0.055	0.045
2	0.215	0.061	0.725	0.060	0.054
3	0.247	0.086	0.667	0.064	0.064
4	0.279	0.111	0.610	0.069	0.075
5	0.312	0.136	0.552	0.074	0.087
6	0.344	0.161	0.495	0.078	0.099
7	0.376	0.186	0.437	0.083	0.111
8	0.408	0.212	0.380	0.088	0.123
9	0.441	0.237	0.323	0.092	0.135
10	0.473	0.262	0.265	0.097	0.148
11	0.505	0.287	0.208	0.102	0.160
12	0.538	0.312	0.150	0.107	0.173
13	0.570	0.337	0.093	0.111	0.185
14	0.602	0.363	0.035	0.116	0.198

Regime 0: The Good Regime

Fund Number	VLACX weight (Large-Cap)	VSMAX weight (Small-Cap)	VBMFX weight (Bonds)	Mean	Standard Deviation
0	0.238	0.004	0.759	0.050	0.036
1	0.276	0.018	0.706	0.055	0.040
2	0.314	0.033	0.653	0.060	0.046
3	0.352	0.048	0.600	0.064	0.052
4	0.390	0.062	0.547	0.069	0.058
5	0.429	0.077	0.494	0.074	0.065
6	0.467	0.092	0.442	0.078	0.072
7	0.505	0.106	0.389	0.083	0.080
8	0.543	0.121	0.336	0.088	0.087
9	0.581	0.135	0.283	0.092	0.094
10	0.620	0.150	0.230	0.097	0.102
11	0.658	0.165	0.178	0.102	0.110
12	0.696	0.179	0.125	0.107	0.117
13	0.734	0.194	0.072	0.111	0.125
14	0.772	0.209	0.019	0.116	0.133

Regime 1: The Bad Regime

Fund Number	VLACX weight (Large-Cap)	VSMAX weight (Small-Cap)	VBMFX weight (Bonds)	Mean	Standard Deviation
0	-0.164	-0.208	1.372	0.050	0.169
1	-0.106	-0.164	1.271	0.055	0.133
2	-0.049	-0.121	1.170	0.060	0.098
3	0.008	-0.077	1.069	0.064	0.066
4	0.066	-0.033	0.968	0.069	0.044
5	0.123	0.010	0.867	0.074	0.050
6	0.180	0.054	0.766	0.078	0.077
7	0.238	0.097	0.665	0.083	0.110
8	0.295	0.141	0.564	0.088	0.146
9	0.352	0.185	0.463	0.092	0.182
10	0.410	0.228	0.362	0.097	0.219
11	0.467	0.272	0.261	0.102	0.257
12	0.524	0.316	0.160	0.107	0.294
13	0.582	0.359	0.059	0.111	0.332
14	0.639	0.403	-0.042	0.116	0.369

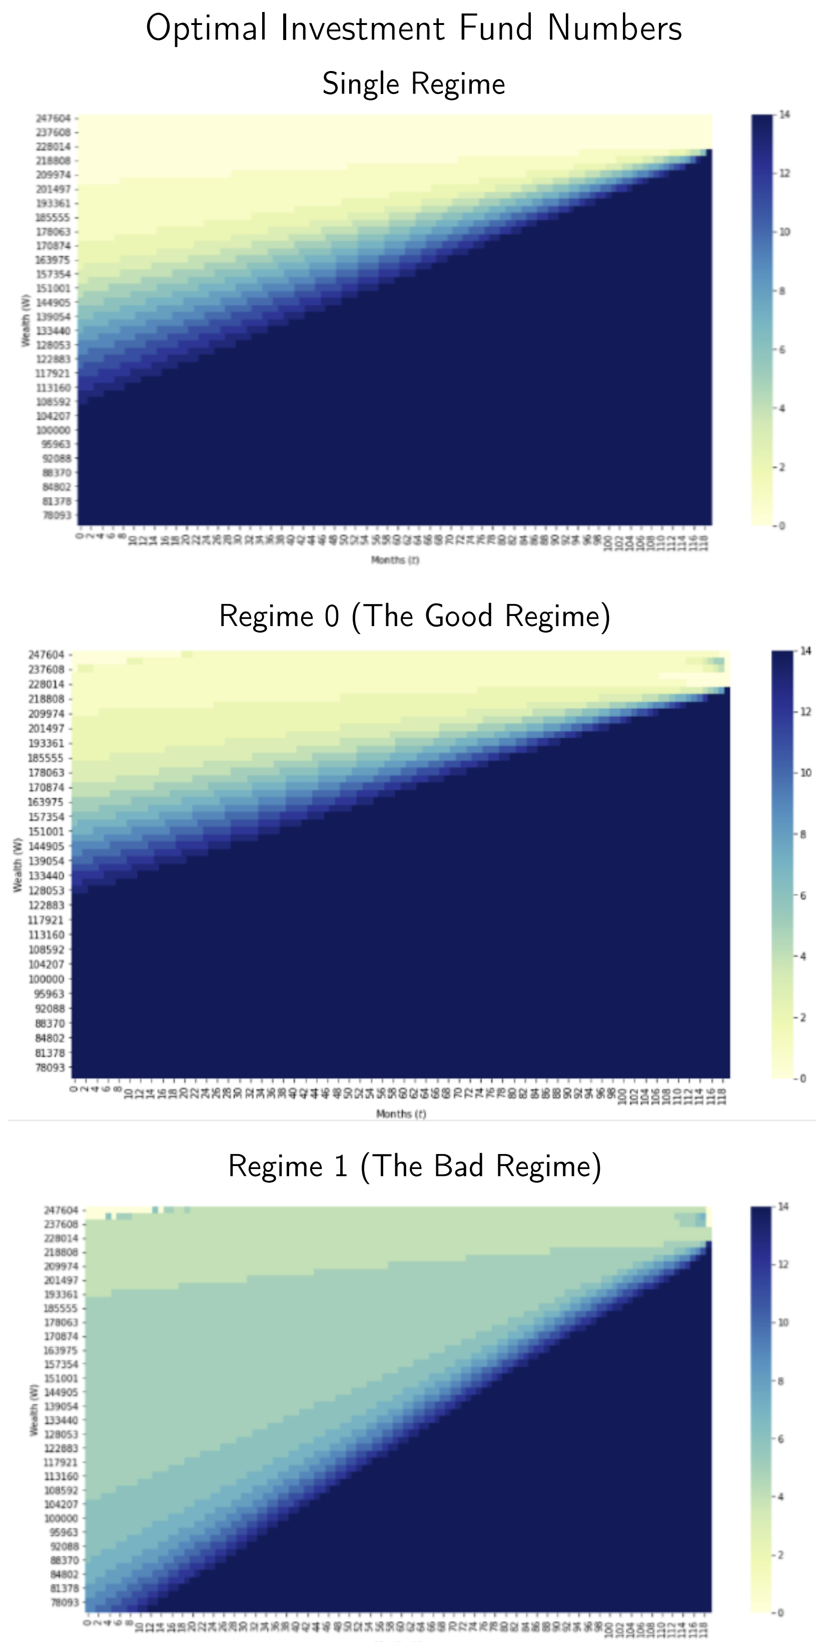

Exhibit 6: Optimal fund choice based on the portfolio's wealth and remaining goal horizon. The top panel shows the optimal fund choice if there is a single regime. The middle and bottom panels show the optimal choice if there are two regimes, depending on the regime. There are $M = 15$ funds in each regime, ranging from the lowest risk fund #0, denoted by the lightest color, to the riskiest, fund #14, denoted by the darkest color. Note that the single regime's results are between the good and bad regimes' results.

4.2 Three regimes versus two regimes

In order to investigate if there is sufficient statistical support to consider a three-regime world, we used the returns for the balanced index fund (VBINX) to estimate the regime-switching model for three regimes. We then compared the two-regime model with the three-regime model using the AIC, BIC, and HQIC metrics, all of which balance goodness-of-fit with the number of regimes employed. For all three metrics, the two-regime model shows lower scores (AIC=966, BIC=986, HQIC=974) than the three-regime model (AIC=970, BIC=1010, HQIC=986). Since the model with the lower scores is better, all three of these metrics suggest that the two-regime model is more appropriate than the three-regime model. In addition, both estimated transition probabilities in the two-regime model are statistically significant (that is, they have p-values < 0.05), whereas only two of the six transition probabilities in the three-regime model are significant, again suggesting that the data does not support extending the model to three regimes.

For the single regime model, the three metrics are AIC=1058, BIC=1065, and HQIC=1061. Since the corresponding numbers for the two-regime model are lower, the use of the two-regime model is, again, justified. However, since there is uncertainty about which of the two regimes is controlling the market, the investor may be better off with a one-regime model instead of a two-regime model. We explore this in the next subsection.

4.3 Comparing optimal goal probabilities for two regimes versus one regime when we have regime certainty

Applying the two-regime dynamic programming model from Appendix A.2 to the example detailed in Subsection 4.1, we can compare the optimal fund strategy for two regimes under the conditions of regime certainty to the optimal fund strategy for one regime.

Suppose there is only one regime and the investor is aware that there is only one regime. In this case, by definition, the investor has perfect knowledge of the regime, and we can solve for the optimal strategy via dynamic programming, using a simplification of the method in Appendix A.2 for a single regime. Applying this to the funds available for the single regime given in the top panel of Exhibit 5, we determine our optimal strategy, which is given by the time and state-dependent optimal fund choices shown in the top panel of Exhibit 6.

How does this compare to the case covered directly in Appendix A.2 where there are two regimes and the investor is aware not only that there are two regimes, but also which regime is governing the market at each time? Because this two-regime model with regime certainty can only benefit from its additional degree of freedom over the single-regime model, it must produce a higher optimal goal probability. Comparing the first and second rows in Exhibit 7, we see that the optimal two-regime goal probability of 0.806 drops to 0.739 for a single regime. This clearly shows that under perfect information, an investor is materially better off having and using two regimes than having and using one regime.

Exhibit 7: Optimal goal probabilities under different conditions. The goal is to grow an initial wealth of \$100,000 to the goal wealth of \$225,000 or more over a horizon of 10 years, with monthly rebalancing. The conditions are given in the first four columns: (1) the actual number of governing regimes, (2) the investor's model for the number of governing regimes, (3) α, the weight given to external information regarding the governing regime, as opposed to $(1-\alpha)$, the weight given to our internal Bayesian model, and (4) β, the probability that the external information will be correct in any given month. The last five columns report the optimal goal probabilities under varying strategies and circumstances. The "Appendix A" column reports the goal probability using the model and information described in Section 3 and Appendix A. The next three columns describe the effect of alternative approaches to using internal information, as described in Subsections 4.4.2–4.4.4. Finally, the "With Inflows and Outflows" column shows the optimal goal probability using the Appendix A method when we add a monthly infusion of \$1000 into the portfolio for the first 60 months and a monthly withdrawal from the portfolio of \$2000 in the remaining 60 months.

Actual Number of Regimes	Investor's Guess for Number of Regimes	α	β	Optimal Goal Probabilities: Appendix A (Subsections 4.3 & 4.4.1)	Changed Cutoff (4.4.2)	Add Single Regime (4.4.3)	Averaging Funds (4.4.4)	With Inflows and Outflows (4.5)
2	2	N/A	N/A	0.806	0.806	0.806	0.806	0.709
1	1	N/A	N/A	0.739	0.739	0.739	0.739	0.648
2	1	N/A	N/A	0.743	0.743	0.743	0.743	0.647
1	2	N/A	N/A	0.716	0.716	0.716	0.736	0.628
2	2	0	0.5	0.674	0.729	0.746	0.700	0.589
2	2	0.2	0.5	0.672	0.724	0.743	0.703	0.585
2	2	0.4	0.5	0.665	0.719	0.730	0.703	0.584
2	2	0.6	0.5	0.665	0.684	0.690	0.696	0.584
2	2	0	0.75	0.674	0.729	0.746	0.700	0.589
2	2	0.2	0.75	0.714	0.737	0.750	0.733	0.623
2	2	0.4	0.75	0.732	0.739	0.749	0.747	0.635
2	2	0.6	0.75	0.732	0.734	0.740	0.746	0.635
2	2	0	1	0.674	0.729	0.746	0.700	0.589
2	2	0.2	1	0.765	0.747	0.757	0.765	0.663
2	2	0.4	1	0.804	0.760	0.770	0.799	0.705
2	2	0.6	1	0.804	0.799	0.802	0.806	0.705

Two other questions naturally arise: (1) What if the markets are governed by two regimes, but the investor mistakenly believes there is only a single regime and therefore uses the fund strategy that is optimized for a single regime? In this case, the investor achieves an optimal goal probability of 0.743 (third row of Exhibit 7), which is close to the optimal goal probability if there actually was only one regime. (2) Next is the opposite question: What if the markets are governed by a single regime, but the investor mistakenly believes there are two regimes and therefore uses the fund strategy that is optimized for two regimes from Subsection A.3? In this case, the investor achieves an optimal goal probability of 0.716 (fourth row of Exhibit 7), which, again, is close to optimal goal probability if they had known the market was governed by a single regime.

Therefore, in this case we see that an investor is not greatly harmed by a lack of awareness of regimes. However, an investor who is cognizant of regimes and has regime certainty, will do better using a two regime model. But even if we assume the market is governed by two regimes, does the two regime model continue to be superior if the investor lives in the real world where there is regime uncertainty?

4.4 The effect of regime uncertainty and external regime information

Assume the market is controlled by two regimes and the investor is looking to optimize their strategy using a two regime model, but there is regime uncertainty. Because the investor does not know when each regime is governing the market as they did with the dynamic programming method in Appendix A.2, they instead use the internal information method for regime uncertainty detailed in Appendix A.3 to guess the governing regime via Bayesian updating.

The resulting optimal goal probability when we have regime uncertainty is 0.674, a considerable decrease from the optimal goal probability of 0.806 when we have regime certainty. Indeed, as we have just seen, only using the single-regime strategy yields better results, with an optimal goal probability of 0.743.

We can also consider another simple strategy: since we know for our example that, on average, we are in Regime 0, the good regime, 3/4 of the time, we can choose a strategy where the investor simply assumes they are always in the good regime, that is, the guess for the governing regime, $R(t)$, is $R(t) = r_0$ for all t, knowing that this will be right about 75% of the time. With this simple strategy, the optimal goal probability becomes 0.717, which is still much higher than 0.674. This is not surprising since the probability of getting the regime correct is 75% here, whereas the internal Bayesian method correctly determines the regime only 68.5% of the time in our example.

4.4.1 External information

The estimate for determining which regime is governing the markets may be improved by external information, meaning we can augment the internal Bayesian update method in Appendix A.3 by using the method detailed in Appendix A.4 for including external information. In these subsections, we set a parameter $\alpha \in (0, 1)$ to decide how much weight to give to information outside the Bayesian update model, which means the information inside the Bayesian update is given a weight of $(1 - \alpha)$. Therefore, if $\alpha = 1$, the investor completely relies on their external information, and if $\alpha = 0$, they completely discard it.

Of course, the external information may be poor. In Appendix A.4 we model this aspect with $\beta \in (0.5, 1)$, the probability of the external information being correct. For example, if $\beta = 1$, the external information is always correct, and if $\beta = 0.5$, the external information has an equal chance of guessing the regime incorrectly as it does correctly, like using a coin toss. Our technical framework allows the investor to choose various levels of accuracy of external information through the parameter $\beta \in (0.5, 1)$, and then they may check if their external regime detection model has a β level that justifies the use of the approach.

The analytic framework used here is agnostic to the external regime detection algorithm used. For example, Resnick and Shoesmith (2002) presents a good example of how external information (specifically, the yield curve shape) may be used to detect regimes in the S&P 500 index. An investor employing this method may check if it yields a β level that justifies this approach to using external information. The Appendix A column in Exhibit 7 quantifiably shows that if you have a good external model to determine which regime is governing the market (i.e, a sufficiently high β for your chosen α), then indeed, a two-regime model is recommended.

Unsurprisingly, using a coin toss to guess the regime decreases the optimal goal probability, which is reflected in Exhibit 7 by the decrease in the optimal goal probability as α increases in the rows where $\beta = 0.5$. Similarly, in the rows where $\beta = 1$, we see an increase in the optimal goal probability as α increases. We note that by $\alpha = 0.2$, our optimal goal probability, 0.765, has exceeded the approximate optimal goal probabilities attained by previous strategies for regime uncertainty. When $\beta = 1$ and $\alpha = 1$, we are back to regime certainty. The small difference in the table between the $\alpha = 1, \beta = 1$ optimal goal probability, which is 0.804, and the value determined by dynamic programming, 0.806, is explained by the fact that 0.804 is an approximation calculated from 10,000 simulated paths, as discussed in Appendix A.3.

Viewing the $\beta = 0.75$ rows in Exhibit 7, we still see that the optimal goal probability increases as α increases. This is no surprise since 0.75 is greater than 0.685, the fraction of times the internal Bayesian method guesses the governing regime correctly. What if we set $\beta = 0.685$, so both the internal and external estimates for the governing regime are correct approximately equally often? Unsurprisingly, in this case the optimal goal probability is changed very little by changing α. It goes up a little from 0.674 when $\alpha = 0$ to 0.707 when $\alpha = 0.2$ and then to 0.718 when $\alpha \geq 0.4$.

The analysis summarized by the Appendix A column in Exhibit 7 sparks an interesting question: How correct does a regime-detection model need to be to yield benefits to goals-based investors? Every practitioner has a different approach to using external information to estimate which regime they believe is governing the market. We aim to give practitioners a tool to help determine whether their approach is accurate enough to be useful. The key question for practitioners here is: if you rely solely on your external information model for predicting the regime governing the market, what minimum value must β have for the two-regime approach with this model to yield a higher probability of achieving the goal wealth than 0.743, which is the probability if you employ the one-regime model? The answer is obtained by setting $\alpha = 1$ and running the model for various values of β until we find the cutoff at which the two-regime model's goal probability exceeds 0.743.[2] It turns out that for the data we use, this cutoff value is $\beta = 0.79$. For the case with infusions and withdrawals (in the last column of Exhibit 7), the cutoff (where the goal probability exceeds 0.647) is $\beta = 0.80$, a similar value.

4.4.2 Alternative Bayesian internal regime information approach: Change the probability cutoff for choosing the regime

Can we use our internal information more effectively? In each period we update the Bayesian probabilities of being in the good and bad regime. In Appendix A.3, we assume that the operative regime is the one with the higher probability. That is, we choose $R(t) = r_0$ if $P_t(r_0) \geq 0.5$; otherwise we choose $R(t) = r_1$. This implicitly assumes a regime cutoff probability of 50%. However, noting that the good regime spans 3/4 of the historical record, perhaps it makes more sense to instead choose a regime cutoff probability of 75%, meaning $R(t) = r_0$ if $P_t(r_0) \geq 1 - 0.75 = 0.25$; otherwise we choose $R(t) = r_1$.

As it turns out, it does make more sense to do this, as the numbers in the "Changed Cutoff" column in Exhibit 7 show. One may wonder if a different regime cutoff probability would be even better. We found, however, that the probabilities of achieving the goal generally decreased, both as we increased the cutoff probability to 0.8 and then 0.9 and when we decreased the regime cutoff probability to 0.6 and then 0.5, it's previously assumed value.

4.4.3 Alternative Bayesian internal regime information approach: Use the single regime strategy when less certain of the governing regime

When $P_t(r_0)$ is close to 1, we naturally wish to choose $R(t) = r_0$, and when $P_t(r_0)$ is close to 0, we naturally wish to choose $R(t) = r_1$. Up until now, we have chosen either $R(t) = r_0$ or r_1 at intermediate values of $P_t(r_0)$ to determine our fund, but what if we instead choose our fund from the single-regime results at these intermediate values?

Of course, there is a question as to what range is chosen to be "intermediate values." We considered three ranges centered at 0.25 = 1 - 0.75 (the optimal regime probability cutoff), specifically the ranges (0.2, 0.3), (0.15, 0.35), and (0.10, 0.40), and we also considered two ranges

[2]Actually for the approach in Appendix A, there is no difference in the results if $\alpha = 0.5$, $\alpha = 1.0$, or α equals any value between. In each of these cases, the regime selected by the external information will be assumed to be correct and the internal information is irrelevant.

centered at 0.5, specifically the ranges (0.25, 0.75) and (0.10, 0.90). Of these five ranges, we found that the range (0.25, 0.75) generally yielded the best results, and these results are shown in the "Add Single Regime" column in Exhibit 7. We note that in all cases, this approach provides a modest improvement over the "Changed Cutoff" column in Exhibit 7.

4.4.4 Alternative Bayesian internal regime information approach: Averaging fund choice

In place of using the single-regime results, the investor could instead hedge their regime bet by using a fund that is a mixture of the optimal funds for regimes 0 and 1, weighted by $P_t(r_0)$ and $P_t(r_1) = 1 - P_t(r_0)$ respectively. We show the effect of this weighted average strategy in the "Averaging Funds" column in Exhibit 7. We note that this strategy does notably worse than the "Add Single Regime" column in Exhibit 7 when α is small, but then becomes a little better than the "Add Single Regime" column as α grows. This method also has the benefit of not requiring a determination of arbitrary parameters.

We see that all three of our alternative Bayesian approaches yield superior results to our basic method in Appendix A.3, except when both α and β are high. When α and β are high, our knowledge and use of the correct regime is high, which means that hedging our regime estimate through alternative approaches is no longer advantageous.

4.5 The effect of portfolio infusions and withdrawals

The "With Inflows and Outflows" column in Exhibit 7 shows results for the case when the first half of the 10-year strategy includes infusions of $1000 each month and the second half has monthly withdrawals of $2000. The fact that we withdraw more money than we put in reduces each of the optimal goal probabilities in this column by around 10 percentage points (under regime certainty, and the difference is smaller when there is uncertainty) from their corresponding values in the other columns, which have no infusions or withdrawals. That is, the infusions and withdrawals do not change the qualitative nature of the results, and, more specifically, the two-regime model outperforms the one-regime model for a similar range of α and β values.

4.6 Robustness of the efficient frontier and regime probabilities

As stated at the beginning of this section, we have explored the effects of regime uncertainty in the sense of not being certain which of the two regimes we are in, but assuming that we know the efficient frontiers of both regimes. As discussed in Subsection 3.1, our efficient frontiers come from considering monthly returns over the ~ 17 year period from 2/1/2004 to 3/31/2021. We use the returns in this period from a balanced index fund asset (VBINX) to determine the two regimes. We then determine the efficient frontier from the returns within each regime for a Large-Cap fund (VLACX), a Small-Cap fund (VSMAX), and a bond fund (VBMFX).

We make the following observations based on experiments we ran on subperiod data, shown in Appendix B. First, the efficient frontiers for the first half of our sample period are very different from that of the second half, so if we estimated the two-regime model from the first half and used it in the second half the error would be considerable. This suggests that two-regime models must have good forward estimation of regime-based efficient frontiers. But good capital market expectations are also important for single-regime models.

Second, the calculated probability of being in a specific regime at a given historical time can also vary depending on the subperiod chosen to be partitioned into two regimes. A comparison of estimated regime probabilities using the whole period of return data versus using just the first subperiod lead to nearly identical results during the first subperiod. But this is not the case for the second subperiod, where the full period's return data gives much higher probabilities of being in the good regime than the second subperiod's return data. In other words, estimating regime probabilities for any period of time (even though we are looking back) is dependent on the time spanned by the time series used. Even for two almost overlapping time series where we only add three additional years to our ~ 17 years of data, the estimated regime probabilities and efficient frontiers can be very different. Estimation risk is always high, and may be higher for two-regime models because more parameters (regime probabilities, two frontiers) need to be estimated. Again, see Appendix B for more details.

5 Conclusions

This paper combines dynamic programming and Bayesian techniques to determine the optimal dynamic investor strategy for attaining a goal wealth when the market is subject to regime switching. Implementing and deploying a portfolio strategy in a regime-switching context entails considering three sources of uncertainty: (i) stochastic returns; (ii) uncertainty about which regime is governing the market at each point in time; and (iii) estimation risk, i.e., uncertainty about the nature of each regime's probability distribution for returns, since a distribution based on data from a past time period may not apply particularly well to a future period. We note that (i) and (iii) are also present in single regime models, which tend to be robust given that fewer parameters are involved. Uncertainty (ii) is unique to regime-switching models, and we find that unless this uncertainty is managed with a significant level of accuracy, moving from a one-regime allocation strategy to a two-regime allocation strategy may be too risky in a goals-based wealth management context.

We investigated the effect of uncertainty (i) and uncertainty (ii) using return data from February 2004 to March 2021. Applying the Hamilton (1989) regime-switching model to return data for a balanced (60% stock, 40% bond) index fund shows that the market behavior during this time period is best described by a two regime model. We then use the return data to estimate a regime transition matrix between these two regimes. For three index-based assets representing U.S. Large

Cap stocks, U.S. Small Cap stocks, and U.S. bonds, we determine separate efficient frontiers for each of the two regimes. This defines the return distribution for each of 15 possible funds we consider along each regime's frontier. The dynamic programming algorithm then determines which fund is optimal as a function of three state variables: the wealth, the remaining time to the portfolio horizon, and the governing market regime. At any point in time, the first two of these state variables are observable, but the third variable, the governing regime, is not, corresponding to uncertainty (ii). We therefore look to minimize uncertainty (ii) by estimating the governing regime using simple Bayesian updating procedures that work forwards in time, unlike dynamic programming, which works backwards in time. We also allow for augmenting these Bayesian approaches with external information about the governing regime, where we specify a parameter, $\alpha \in [0, 1]$, for the weight given to this external information and another parameter, $\beta \in [0.5, 1]$, for the probability that the regime specified by the external information is correct.

We find that this two-regime strategy can produce results for a goals-based investor that are superior to a one-regime strategy (that is, a strategy that ignores regimes), but only when there is a sufficiently good idea of the governing regime, i.e., when uncertainty (ii) has a limited effect. Otherwise, even if investors correctly know there are two governing regimes, their uncertainty about which regime is governing the market can make them less likely to attain their goal than an investor who ignores regimes. We provide a metric that allows us to quantify how good knowledge of the governing regime must be to prefer the two-regime model. Further, for the two-regime model, we explore alternative Bayesian methods that hedge regime forecasts, and we quantify how effective these alternative approaches are as partial antidotes to regime uncertainty.

When a goal-driven portfolio is underperforming, our results reflect the optimal strategy of taking on additional risk to clear the performance backlog, in contrast to the typical investor behavior discussed in the introduction. However, it does this in a measured way, taking less risk when the governing regime has a poorer risk-return tradeoff. The investment strategy in this paper offers a more precise way to navigate the regime changes we experience and also suggests well-tuned prescriptions that fit with the strategy of being more aggressive with investments when the market goes down. Therefore, our algorithm may be used to navigate transitions between bull and bear markets, but also warns that the success of such a strategy depends materially on good forecasts for which regime is governing the markets as time evolves. From our data, we conclude that pure Bayesian updating ($\alpha = 0$) is not sufficient to build a dynamic investment strategy for managing wealth across regimes to obtain a goal. We also need good external information ($\beta \to 1$) to make a two-regime algorithm more successful than the standard one-regime model for goals-based investors.

An important extension of this paper is to examine the role of hedging, see Bhansali (2013). That is, when assuming we are in the bad regime, instead of transitioning the good-regime fund to the bad-regime fund, the portfolio is retained as is, and downside hedges are added instead, possibly using derivatives. Without a detailed analysis, it is hard to speculate whether the hedging approach will do better and under what conditions. A blended approach of using GBWM with options is needed for this, as has been explored in Das and Ross (2020), and such an approach, combined with regime-switching, will make for an interesting technical and empirical analysis.

A Appendix: Technical Details of the Two-Regime Goal Probability Optimization Algorithm

A.1 The wealth transition probabilities

We establish some notation before explaining the algorithm we use to solve the problem. We define the state space using a three-dimensional tensor, where the three dimensions correspond to time $t \in \{0, 1, 2, ..., T\}$, wealth $W(t) \in \{W_0, W_1, ..., W_{i_{\max}}\}$, and regime $r(t) \in \{r_0, r_1\}$, where r_0 is the good market regime and r_1 is the bad market regime. The time interval between successive t values in the model is defined to be h, so if $h = 1$ then t is in years, and if $h = \frac{1}{12}$ then t is in months. At each $t = 0, 1, ..., T-1$, there is a possible cash infusion into the portfolio, denoted by $I(t)$, and a possible consumption, denoted by $C(t)$. The net inflow, of course, is $I(t) - C(t)$.

We aim to determine the optimal fund at each point in the state space $\{t, W(t), r(t)\}$. The optimal fund will reside on one of the two regime-specific efficient frontiers, each determined using Markowitz (1952) mean-variance optimization. For programming convenience, we have set the number of funds, M, to be the same arbitrary, but finite, value in each regime. Therefore, within each regime we have an ordered set of funds with annual means $\mu_1 < \mu_2 < ... < \mu_M$ and, from the regime's efficient frontier, corresponding annual standard deviations $\sigma_1 < \sigma_2 < ... < \sigma_M$. We will use the index $k = 1, 2, ..., M$ to denote a specific fund, so fund k has mean $\mu_{k(r_l)}$ and standard deviation $\sigma_{k(r_l)}$, where l, the index of the regime, is either 0 or 1.

The portfolio's wealth evolves from its known initial value, $W(0)$, under the governance of a Markovian evolution model. In this paper, we have chosen geometric Brownian motion for our evolution model, so

$$W(t+1) = [W(t) + I(t) - C(t)]e^{\left(\mu - \frac{1}{2}\sigma^2\right)h + \sigma\sqrt{h}Z}, \tag{1}$$

where Z is a standard normal variate. This Gaussian random variable may be replaced (without loss of generality) with other types of stochastic terms, for example, we could draw from a T-distribution instead of the Z-distribution if we wanted fatter tails. Rearranging equation (1) for a given fund, $k(r_l)$, allows us to define the following measure of the likelihood that a portfolio at wealth level $W_i(t)$ will transition to the wealth level $W_j(t+1)$:

$$P(W_j(t+1)|W_i(t), k(r_l)) = \phi\left[\frac{1}{\sigma\sqrt{h}}\left(\ln\left(\frac{W_j(t+1)}{W_i(t) + I(t) - C(t)}\right) - \left(\mu_{k(r_l)} - \frac{1}{2}\sigma^2_{k(r_l)}\right)h\right)\right],$$

where $\phi[\cdot]$ is the standard normal probability density function. We then determine Q, the wealth transition probabilities from wealth level $W_i(t)$ to any wealth level $W_j(t+1)$, by normalizing P:

$$Q(W_j(t+1)|W_i(t), k(r_l)) = \frac{P(W_j(t+1)|W_i(t), k(r_l))}{\sum_{m=0}^{i_{\max}} P(W_m(t+1)|W_i(t), k(r_l))},$$

which gives us that $\sum_{j=0}^{i_{\max}} Q(W_j(t+1)|W_i(t), k(r_l)) = 1$; that is, the probabilities add to one, as they must.

For simplicity, the wealth grid, $\{W_0, W_1, ..., W_{i_{\max}}\}$ is independent of t. In order to ensure the grid is appropriately dense, we space the grid so that the logarithm of the wealth grid's values are equidistant (therefore the wealth grid values themselves are not equidistant), which better reflects the structure of equation (1). The value of $W_{i_{\max}}$ is determined from: (1) considering the effects of $I(t)$ and $C(t)$, (2) using the larger of the two $\mu_{M(r_l)}$ (where $l = 0$ or 1), and (3) using returns that correspond to being three standard deviations above the mean, where the standard deviation is the larger of the two $\sigma_{M(r_l)}$. The same method, but using the smaller of the two $\mu_{M(r_l)}$ and returns three standard deviations below the mean are used to determine the smallest grid point, W_0.

A.2 Backward recursion with dynamic programming and the Bellman equation

We wish to determine the value function $V(t, W_i, r_l)$, which represents the maximum probability that we will have attained a goal wealth, G, at the final time T, given that at time t we are in regime r_l and the portfolio wealth is W_i. The value function is determined via backward recursion in time on our three-dimensional state space. At each triple, $\{t, W_i, r_l\}$, the action space is a choice among one of the M funds in the extant regime, r_l.

At the portfolio horizon T, the value function only depends on whether or not the goal is met. That is,

$$V[T, W_i(T), r_l(T)] = \begin{cases} 1 & \text{if } W_i(T) \geq G \\ 0 & \text{if } W_i(T) < G. \end{cases}$$

With the value function at time $t = T$ in place, we then compute the value function at time $t = T-1$, then $t = T-2, ...$, then finally $t = 0$, by repeatedly applying the Bellman (1952) equation

$$\begin{aligned} V[t, W_i, r_l] \quad &= \quad \max_{k(r_l)} \sum_{j=0}^{i_{\max}} \Big[V[t+1, W_j, r_l] \cdot Q(W_j|W_i, k(r_l)) \cdot p_{ll} \\ &+ \quad V[t+1, W_j, r_{(1-l)}] \cdot Q(W_j|W_i, k(r_l)) \cdot p_{l(1-l)} \Big], \end{aligned}$$

where p_{00}, p_{01}, p_{10}, and p_{11} are the elements of $\mathbf{p}$, the 2×2 regime transition probability matrix from Subsection 3.1. At each node (t, W_i, r_l) in the state space, the optimal fund $k \in 1, 2, ..., M$

is the fund that maximizes the right-hand side of the Bellman equation. We store this fund in a separate 3D tensor, denoted $K(t, W_i, r_l)$. This allows us access to the optimal fund at any node in the state space.

An important caveat in the Bellman equation is that the optimal fund is chosen assuming that the investor has *regime certainty*; that is, the investor always knows which regime they are in, so the triple, $\{t, W_i, r_l\}$, is known. In practice, as time moves forward, t and $W_i(t)$ are, of course, known, but $r_l(t)$ is unobservable and needs to be estimated. This *regime uncertainty* means the attainable probability of achieving the goal wealth G at time T is less than the probability given by the value function. The better the estimate of the regime $r_l(t)$ is, the smaller the difference between the attained probability and the optimal probability given by the value function. In the next subsection we show how to use Monte Carlo simulation, which works forwards in time, to determine a Bayesian estimate for the governing regime at each time.

A.3 Forward in time estimation of the regime

A critical question we explore in this paper is whether an investor is better off using a model for a single regime or for two regimes, even if they know returns are drawn from two regimes. If the investor has regime certainty, they are guaranteed to be better off with the two regime model, but the effects of regime uncertainty in the two regime model may be severe. That is, if the investor is unable to use the information at hand to determine which regime they are in, they may end up holding the wrong fund much of the time, making them worse off than naively assuming a single regime, where this risk is not present. In the hopes of reducing the ill effects of regime uncertainty, this subsection uses Monte Carlo simulation, the information available at each point in time, and Bayesian methods to estimate which regime is operative. The next subsection will show how to incorporate this subsection's use of information that is internal to our model with information that is external to our model to estimate the governing regime. After that, in Section 4, we will be able to explore the answer to our critical question by presenting results from experiments with markets governed by two regimes and then seeing if the goal probabilities obtained from the one-regime model or the two-regime model are better.

We start our simulation at time $t = 0$, where $W(0) = W_{i_0}$ is the initial wealth, $r(0)$ is the true economic regime, and $R(0)$ is the investor's guess for the economic regime. The investor uses the K-tensor, specifically $K(0, W_{i_0}, R(0))$, to determine their fund, which they believe to be $K(0, W_{i_0}, r(0))$, the optimal fund choice if they had regime certainty.

At each time t, the investor's fund choice, $K(t, W_i(t), R(t))$, determines the weights of the three asset classes, which gives annualized values of μ and σ in the true regime, $r(t)$. Equation (1) uses these values of μ and σ, along with $h = \frac{1}{12}$ for monthly returns, to generate a new value of wealth at time $t + 1$, which corresponds to the next month. The investor will observe the realized return of the three asset classes, and use it to help determine $R(t + 1)$, their guess at time $t + 1$ of the true regime, $r(t + 1)$. This is accomplished using a Bayesian updating process, as we show below. We repeat this process for $t = 0, 1, ..., T$ to complete one path in the forward simulation, recording at the end whether or not the strategy met its goal of attaining a wealth of at least G at time T. This procedure from $t = 0$ to $t = T$ is repeated for 10,000 simulated paths, with the fraction of paths that attain the goal G generating a close approximation of the probability of attaining the goal wealth G.

We now describe the details of the Bayesian updating process as we move from a given time t to time $t + 1$ on a single path:

1. At time t, we define $P_t(r_0)$ and $P_t(r_1)$ to be our previously estimated probabilities of being in each of the two regimes. (At $t = 0$, since there is no previous estimate, we define $P_0(r_0) = \frac{\frac{1}{p_{01}}}{\frac{1}{p_{10}} + \frac{1}{p_{01}}}$ and $P_0(r_1) = \frac{\frac{1}{p_{10}}}{\frac{1}{p_{10}} + \frac{1}{p_{01}}}$, which are the average probabilities over time of being in each regime.) If $P_t(r_0) \geq P_t(r_1)$ (or, equivalently, $P_t(r_0) \geq 0.5$), then the investor assumes they are in Regime 0 at time t, that is, $R(t) = r_0$. In this case they select fund $K(r_0)$, which is an abbreviation for $K(t, W_i(t), r_0)$ where $W_i(t)$ is the closest grid point to the actual wealth, $W(t)$. If $P_t(r_0) < 0.5$, then the investor selects $K(r_1)$.

2. Of course it is possible that $R(t)$, the guessed regime, is not equal to $r(t)$, the true regime used to determine the asset returns between time t and time $t + 1$. We must therefore consider the likelihood of the observed new wealth $W(t + 1)$ if the true regime, $r(t)$, is either r_0 or r_1. Recalling that $\phi(\cdot)$ is the probability density function for the standard normal distribution, this likelihood is defined by

$$L(W(t+1)|r_l) = \phi\left(\frac{\ln\left(\frac{W(t+1)}{W(t)+I(t)-C(t)}\right) - \left(\mu(K(R(t)), r_l) - \frac{(\sigma(K(R(t)), r_l))^2}{2}\right)h}{\sigma_i\sqrt{h}}\right),$$

   where $\mu(K(R(t)), r_l)$ and $\sigma(K(R(t)), r_l)$ are the annual mean and standard deviation of the fund in regime r_l with asset weights given by $K(R(t))$. So, for example, if $R(t) = r_0$, and fund $K(r_0)$ corresponds to having 50% of the portfolio in the large-cap index (VLACX), 30% in the small-cap index (VSMAX), and 20% in the bond index (VBMFX), but we consider being in regime r_1, then, using the mean returns in Regime 1 given in Exhibit 4 for the three assets, we have that

$$\mu(K(r_0), r_1) = 0.1114 \cdot 0.5 + 0.1183 \cdot 0.3 + 0.0678 \cdot 0.2 = 0.10475.$$

   Note that if $R(t) = r(t) = r_l$, then $\mu(K(R(t)), r_l) = \mu_{K(r_l)}$ and $\sigma(K(R(t)), r_l) = \sigma_{K(r_l)}$.

3. Applying Bayes' theorem to $L(W(t + 1)|r_l)$, which is the prior distribution, and to $P_t(r_0)$ and $P_t(r_1)$, the estimated probabilities for being in each of the two regimes at time t, we

obtain the posterior probability distribution:

$$\begin{aligned}
f_t(r_0|W(t+1)) &= \frac{L(W(t+1)|r_0)\cdot P_t(r_0)}{L(W(t+1)|r_0)\cdot P_t(r_0)+L(W(t+1)|r_1)\cdot P_t(r_1)},\\
f_t(r_1|W(t+1)) &= \frac{L(W(t+1)|r_1)\cdot P_t(r_1)}{L(W(t+1)|r_0)\cdot P_t(r_0)+L(W(t+1)|r_1)\cdot P_t(r_1)}\\
&= 1-f_t(r_0|W(t+1)).
\end{aligned}$$

4. Taking into account the probability of switching regimes given by the elements in the regime transition probability matrix $\mathbf{p}$, we can now calculate $P_{t+1}(r_0)$ and $P_{t+1}(r_1)$, the estimated probabilities for being in each of the two regimes at time $t+1$:

$$\begin{aligned}
P_{t+1}(r_0) &= f_t(r_0|W(t+1))\cdot p_{00}+f_t(r_1|W(t+1))\cdot p_{10},\\
P_{t+1}(r_1) &= f_t(r_0|W(t+1))\cdot p_{01}+f_t(r_1|W(t+1))\cdot p_{11}\\
&= 1-P_{t+1}(r_0).
\end{aligned}$$

   Having attained these probabilities, we determine $R(t+1)$ for the next time period by setting $R(t+1)=r_0$ if $P_{t+1}(r_0)\geq 0.5$ and $R(t+1)=r_1$ if $P_{t+1}(r_0)<0.5$.

5. Having determined the estimated regime, $R(t+1)$, we must also determine the true regime, $r(t+1)$, in the simulated path. This is simple to determine using the elements of $\mathbf{p}$ and u, a single sample from a uniform distribution between 0 and 1. In the case where $r(t)=r_0$, we stay in Regime 0, that is, $r(t+1)=r_0$, if $u\leq p_{00}$, and we switch regimes, that is, $r(t+1)=r_1$, if $u>p_{00}$. Similarly, if $r(t)=r_1$, we stay in Regime 1 at time $t+1$ if $u\leq p_{11}$, and switch if $u>p_{11}$.

A.4 Enhancing regime forecasts with exogenous information

In the process just described, the investor uses their guess for the regime, $R(t)$, and the observed new wealth, $W(t+1)$, to determine their guess for the next period's regime, $R(t+1)$. However, investors often also incorporate other external information and techniques to form their guess for $R(t+1)$, such as historical data, market forecasts, time-series forecasting methods, etc. Investors will weigh this external information against the values for $P_{t+1}(r_0)$ and $P_{t+1}(r_1)$ calculated above, which will now call $P^{\text{int}}_{t+1}(r_0)$ and $P^{\text{int}}_{t+1}(r_0)$, since they are calculated with internal information. We reflect this weighing via the parameter $\alpha\in[0,1]$, which is the weight given to the external information. We also have a parameter $\beta\in[0.5,1]$, which is the probability with which the external information is accurate. (A β value below 0.5 would mean the external model was wrong more often than it was right, and we assume such a model would be quickly replaced or discarded.) In the simulated paths, the regime forecast from the external information is added to the information internal to the model to obtain a revised forecast via the following two equations:

$$
\begin{aligned}
P_{t+1}(r_0) &= \alpha \cdot \left[I_{\{u \le \beta\}} \cdot I_{\{r(t+1)=r_0\}} + I_{\{u > \beta\}} \cdot (1 - I_{\{r(t+1)=r_0\}}) \right] \\
&\quad + \; (1-\alpha) \cdot P^{\mathsf{int}}_{t+1}(r_0)
\end{aligned}
$$

$$
\begin{aligned}
P_{t+1}(r_1) &= \alpha \cdot \left[I_{\{u \le \beta\}} \cdot I_{\{r(t+1)=r_1\}} + I_{\{u > \beta\}} \cdot (1 - I_{\{r(t+1)=r_1\}}) \right] \\
&\quad + \; (1-\alpha) \cdot P^{\mathsf{int}}_{t+1}(r_1) \\
&= 1 - P^{\mathsf{int}}_{t+1}(r_0),
\end{aligned} \tag{2}
$$

where $I_{\{c\}}$, the indicator function, equals 1 if condition c is met and 0 if condition c is not met, and where u, as before, is a single sample from a uniform distribution between 0 and 1. We note that the same u is used for all four of the $I_{\{u \le \beta\}}$ and $I_{\{u > \beta\}}$ expressions in the two equations above, so both $I_{\{u \le \beta\}}$ equal 0 and both $I_{\{u > \beta\}}$ equal 1 or both $I_{\{u \le \beta\}}$ equal 1 and both $I_{\{u > \beta\}}$ equal 0.

We implement this new method for updating the regime probabilities, $P_t(r_0)$ and $P_t(r_1)$, at each time period within each of our simulated paths. As stated earlier, we simulate 10,000 Monte Carlo paths and approximate the goal probability as the fraction of these paths where the goal G is reached or exceeded at the time horizon T. Because of regime uncertainty, this goal probability should be less than the value function obtained in Section A.2 under regime certainty.

Section 4 presents the results obtained from applying the algorithm in this appendix to the example of the 17 years of monthly returns used in Subsection 3.1 and the corresponding parameters obtained there.

B Robustness to parameter estimation

We briefly discuss the sensitivity of our parameters to the choice of the estimation subperiod. We split our 17.5 year spanning period into two subperiods: from 1/1/2003 to 12/31/2011 and from 1/1/2012 to 6/30/2020. Do the efficient frontiers within each subperiod look the same as those for the full period? In Exhibit 8, we show the efficient frontiers in each subperiod, and note that while their averaged behavior conforms to the full period's frontiers, they show considerably different behavior from each other. In the first subperiod (left panel), the regimes are quite distinct from each other, in contrast to the second subperiod (right panel), where the two regimes' frontiers are quite close. Further, in the first subperiod, the good regime's efficient frontier is generally above the bad regime's, indicating superior performance, on average, in the good regime. In the second subperiod, the opposite occurs. The bad regimes' efficient frontier is generally above, indicating superior performance, on average, in the bad regime. We also note the considerable change in the volatility between the two periods. The first subperiod contained both the Great Recession and the end of the dot-com bust, so it is no surprise to see that its volatility is far higher than that seen in the second subperiod.

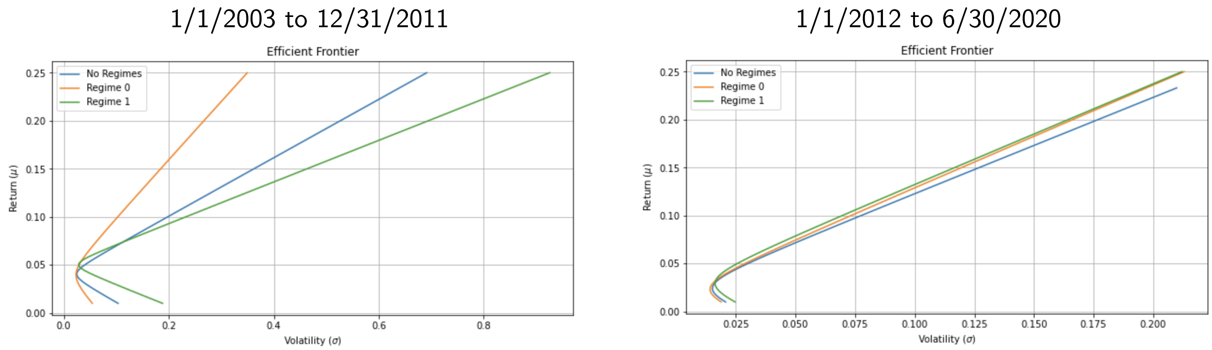

Exhibit 8: Efficient frontiers for two subperiods in the case of a single regime and in the case of two regimes, the good one (Regime 0) and the bad one (Regime 1). The relative behavior of the efficient frontiers within each panel changes considerably between the panels.

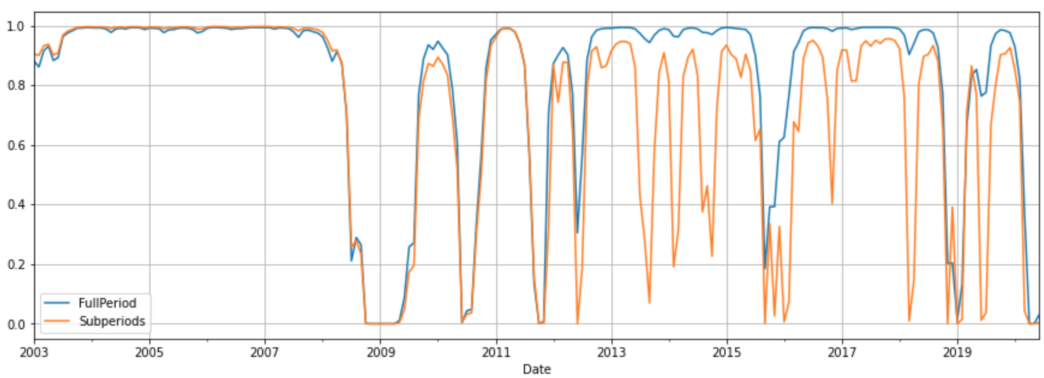

Exhibit 9: Probability of being in the good regime, when estimated using the full sample from 1/1/2003 to 6/30/2020 versus estimates from each sub-period (1/1/2003 to 12/31/2011 & 1/1/2012 to 6/30/2020) spliced together.

As a further assessment of parameter instability, we examined the estimated probability of being in the good regime, first using returns from the entire period from 1/1/2003 to 6/30/2020 and then using just the returns within each of the two subperiods. The comparison plot is shown in Exhibit 9. We see that using the full period of return data versus using the first subperiod of return data lead to nearly identical results during the first subperiod. But this is not the case for the second subperiod, where the full period's return data gives much higher probabilities of being in the good regime than the second subperiod's return data. As with Exhibit 8, this suggests that subdividing input data can lead to very different outcomes.

We can also consider the effect of adding a few years of returns to the data. In Exhibit 10, we compare efficient frontiers from two mostly overlapping periods of return data. The left panel is identical to the efficient frontiers shown in Exhibit 4. That is, we use monthly returns from the 17.5 year period from 1/1/2003 to 6/30/2020. The right panel comes from monthly returns that include three additional years, from 1/1/2000 to 12/31/2002, which includes the majority of the dot-com bust. In the left panel, the good regime's efficient frontier is generally above the bad regime's, meaning the risk-return tradeoffs are better in the good regime. The single-regime efficient frontier is between the good and bad regimes' efficient frontiers.

1/1/2003 to 6/30/2020 | 1/1/2000 to 6/30/2020

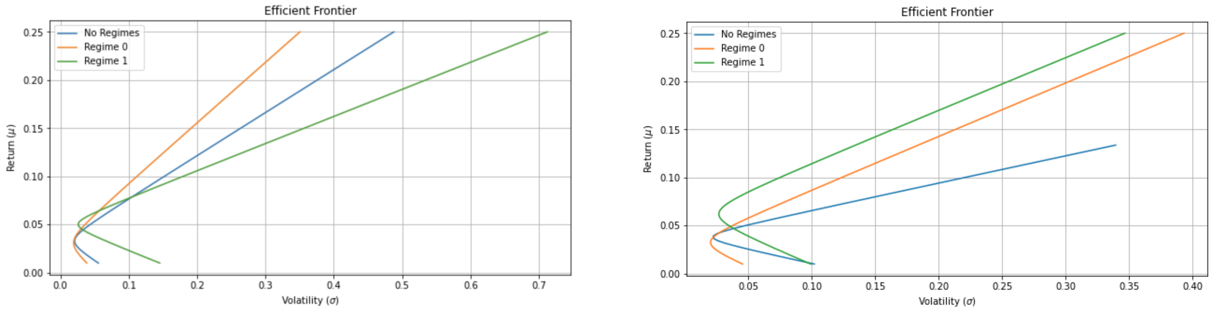

Exhibit 10: Comparison of efficient frontiers determined using slightly different periods of returns. The frontier in the left panel is for the data sample running from 1/1/2003 to 6/30/2020. It is identical to the efficient frontier panel shown in Exhibit 4. The right panel uses data from 1/1/2000 to 6/30/2020, meaning it adds three additional years of data, including most of the dot-com bust. Note the considerable change between these two panels created by the slightly expanded time frame.

The addition of the three additional years of returns, however, changes everything. In the right panel, the efficient frontier for the bad regime is now clearly generally above the frontier for the good regime instead of below it, and the single-regime frontier is now well below both of these regimes' frontiers instead of between them. One may interpret the set of frontiers in the left panel as a world in which growth investors would be better off, whereas in the world described by the right panel, value investing is prescribed since the bad regime is the better regime in which to take additional risk. This suggests that the strategy mentioned in the introduction about taking more risk in a downturn would be even more effective in this goals-based investing paradigm. However, our analysis also suggests that parameter estimation is quite sensitive to the data period chosen, making forecast error an important consideration in choosing a two-regime model versus a one-regime model.

References

Ammann, M. and M. Verhofen (2006, September). The Effect of Market Regimes on Style Allocation. Financial Markets and Portfolio Management 20(3), 309–337.

Ang, A. and G. Bekaert (2002). International Asset Allocation With Regime Shifts. Review of Financial Studies 15(4), 1137–1187.

Ang, A. and G. Bekaert (2004). How Regimes Affect Asset Allocation. Financial Analysts Journal 60(2), 86–99.

Ang, A. and A. Timmermann (2012). Regime Changes and Financial Markets. Annual Review of Financial Economics 4(1), 313–337.

Bellalah, M., A. Hakim, K. Si, and D. Zhang (2020, June). Long term optimal investment with regime switching: inflation, information and short sales. Annals of Operations Research.

Bellman, R. (1952, August). On the Theory of Dynamic Programming. Proceedings of the National Academy of Sciences 38(8), 716–719.

Bernhart, G., S. Höcht, M. Neugebauer, M. Neumann, and R. Zagst (2011, February). Asset correlations in turbulent markets and the impact of different regimes on asset management. Asia-Pacific Journal of Operational Research 28(01), 1–23. Publisher: World Scientific Publishing Co.

Bhansali, V. (2013, December). Tail Risk Hedging: Creating Robust Portfolios for Volatile Markets (1st edition ed.). McGraw-Hill Education.

Branch, W. and G. Evans (2010). Asset Return Dynamics and Learning. Review of Financial Studies 23(4), 1651–1680.

Brunel, J. (2015). Goals-Based Wealth Management: An Integrated and Practical Approach to Changing the Structure of Wealth Advisory Practices. New York: Wiley.

Bulla, J., S. Mergner, I. Bulla, A. Sesboüé, and C. Chesneau (2011, November). Markov-switching asset allocation: Do profitable strategies exist? Journal of Asset Management 12(5), 310–321.

Chhabra, A. B. (2005, January). Beyond Markowitz: A Comprehensive Wealth Allocation Framework for Individual Investors. The Journal of Wealth Management 7(4), 8–34.

Dapena, J. P., J. A. Serur, and J. R. Siri (2019, December). Risk on-Risk off: A regime switching model for active portfolio management. Technical Report 706, Universidad del CEMA. Publication Title: CEMA Working Papers: Serie Documentos de Trabajo.

Das, S. R., D. Ostrov, A. Radhakrishnan, and D. Srivastav (2018). Goals-Based Wealth Management: A New Approach. Journal of Investment Management 16(3), 1–27.

Das, S. R. and G. Ross (2020, September). The Role of Options in Goals-Based Wealth Management. SSRN Scholarly Paper ID 3691699, Social Science Research Network, Rochester, NY.

Das, S. R. and R. Uppal (2004). Systemic Risk and International Portfolio Choice. The Journal of Finance 59(6), 2809–2834.

DeMiguel, V., L. Garlappi, and R. Uppal (2009, May). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? The Review of Financial Studies 22(5), 1915–1953.

Grobys, K. (2012). Active PortofolioManagement in the Presence of Regime Switching: What Are the Benefits of Defensive Asset Allocation Strategies If the Investor Faces Bear Markets? The Review of Finance and Banking 4(1), 15–31. Publisher: EDITURA ASE.

Guidolin, M. and A. Timmermann (2007). Asset allocation under multivariate regime switching. Journal of Economic Dynamics and Control 31(11), 3503–3544.

Guidolin, M. and A. Timmermann (2008). International asset allocation under regime switching, skew, and kurtosis preferences. Review of Financial Studies 21(2), 889–935.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. Econometrica 57(2), 357–384.

Hu, Q. (2019). Essays on asset allocation and delegated portfolio management. Accepted: 2019-10-03T13:42:16Z.

Jiang, P., Q. Liu, and Y. Tse (2015). International Asset Allocation with Regime Switching: Evidence from the ETFs. Asia-Pacific Journal of Financial Studies 44(5), 661–687. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajfs.12109.

Kim, C.-J. and C. Nelson (1999). State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications. MIT Press Books, The MIT Press. ISBN: 9780262112383.

Kupelian, I. (2020, June). Using market regimes, change-points and anomaly detection for Investment Management.

Lewin, M. and C. H. Campani (2020). Portfolio Management under Multiple Regimes: Strategies that Outperform the Market. RAC - Revista de Administração Contemporânea (Journal of Contemporary Administration) 24(4), 300–316. Publisher: ANPAD - Associação Nacional de Pós-Graduação e Pesquisa em Administração.

Markowitz, H. H. (1952). Portfolio Selection. Journal of Finance 6, 77–91.

Merton, R. (1973). An Intertemporal Capital Asset Pricing Model. Econometrica 41(5), 867–87.

Nevins, D. (2004, January). Goals-Based Investing: Integrating Traditional and Behavioral Finance. The Journal of Wealth Management 6(4), 8–23.

Pagan, A. R. and K. A. Sossounov (2003). A Simple Framework for Analysing Bull and Bear Markets. Journal of Applied Econometrics 18(1), 23–46.

Parker, F. J. (2020, April). Allocation of Wealth Both Within and Across Goals: *A Practitioner's Guide*. The Journal of Wealth Management 23(1), 8–21.

Parker, F. J. (2021, January). A Goals-Based Theory of Utility. Journal of Behavioral Finance 22(1), 10–25. Publisher: Routledge _eprint: https://doi.org/10.1080/15427560.2020.1716359.

Resnick, B. G. and G. L. Shoesmith (2002, May). Using the Yield Curve to Time the Stock Market. Financial Analysts Journal 58(3), 82–90.

Shefrin, H. and M. Statman (2000). Behavioral Portfolio Theory. The Journal of Financial and Quantitative Analysis 35(2), 127–151.

Singhal, S. and P. C. Biswal (2019, January). Dynamic Commodity Portfolio Management: A Regime-switching VAR Model. Global Business Review, 0972150918811705. Publisher: SAGE Publications India.

Tu, J. (2010). Is Regime Switching in Stock Returns Important in Portfolio Decisions? Management Science 56(7), 1198–1215.

Vo, H. T. and R. Maurer (2013, March). Dynamic Asset Allocation with Regime Shifts and Long Horizon CVaR-Constraints. SSRN Scholarly Paper ID 2191286, Social Science Research Network, Rochester, NY.

Zhou, X. Y. and G. Yin (2003, January). Markowitz's Mean-Variance Portfolio Selection with Regime Switching: A Continuous-Time Model. SIAM Journal on Control and Optimization 42(4), 1466–1482.