Regime-based Portfolio construction

Ruipeng Deng
Cornell CFEM

Nov 2025

Abstract

In this report, We examine regime-based dynamic asset allocations, compare our method with traditional static 60/40 portfolio in different market environments. We construct a Hidden Markov Model (HMM) using equity returns and yield-curve features to extract four latent market regimes, each displaying distinct and economically interpretable behavior. These states range from risk-on expansions and duration-led flight-to-quality periods to stressed slow-growth environments and acute crisis episodes. Asset-level performance within each regime is analyzed across U.S. equities, Treasuries, investment-grade credit, and high-yield bonds.

We then use identified regimes to guide dynamic portfolio construction under several allocation rules: equal weight, maximum return, maximum Sharpe ratio, minimum variance. For each specification, portfolio weights are conditioned on the prevailing latent state, producing a set of regime-aware strategies. We evaluate their performance relative to the static 60/40 benchmark, focusing on annualized returns, volatility, draw-downs, and risk-adjusted metrics.

The findings illustrate how incorporating probabilistic regime information can reshape portfolio behavior and highlight both the advantages and practical limitations of using state-dependent allocation in modern multi-asset investing.

Contents

1 Introduction

The goal of this paper is to evaluate whether a regime-based dynamic asset allocation strategy can outperform the traditional 60/40 stock–bond portfolio in an environment of low real yields and unstable stock–bond correlations. The 60/40 mix remains a widely used policy benchmark, yet it implicitly assumes a single stable return distribution and a fixed risk relationship between equities and bonds. Recent research on time variation in returns, volatility, and cross-asset correlations shows that these assumptions are fragile. Financial markets cycle through crash, slow-growth, and recovery states with distinct risk–return trade-offs, and the stock–bond correlation frequently changes sign across these environments ([Ang and Bekaert, 2002]; [Guidolin and Timmermann, 2007]). When regimes shift, a static 60/40 allocation can be poorly positioned, increasing drawdowns precisely when downside protection is most needed. This motivates a dynamic approach that adapts to changing market states rather than assuming a constant underlying structure.

1.1 Literature Background

A substantial literature builds portfolio strategies around regime-switching models. Studies using multivariate Markov-switching and hidden Markov models (**HMM**) find that a small number of latent regimes explains much of the joint behavior of stock and bond returns, and that portfolios which adjust risk exposure across these regimes can improve Sharpe ratios and reduce downside risk relative to static benchmarks ([Nystrup et al., 2018]; [Guidolin and Timmermann, 2007]; [Abrahamsen and Nakovski, 2021]). More recent contributions integrate regime forecasts into standard optimization frameworks—such as minimum-variance or mean–variance portfolios—and often use liquid exchange-traded funds to ensure that the strategies are implementable in practice ([Uysal and Mulvey, 2021]; [Bae et al., 2014]). Practitioner-oriented work, including Ilmanen's Expected Returns, emphasizes that simple macro variables such as the slope of the yield curve and broad equity performance carry much of the information investors use when identifying risk-on or risk-off conditions. Our method sits at the intersection of these approaches: it applies regime-switching methods, keeps the signal set intentionally simple, and tests whether a transparent allocation rule can generate superior risk-adjusted performance relative to the static 60/40 benchmark.

1.2 Data and Model

We construct a four-state Gaussian hidden Markov model using two intuitive features: monthly log returns on the **S&P 500** and monthly changes in the **Treasury term spread**, taken from the McCracken FRED–MD dataset. The model is estimated on **post-1990** data and its inferred regimes are mapped onto a four-asset universe representative of standard policy portfolios: U.S. equities (**S&P 500 total return**), intermediate Treasuries (**IEF**), investment-grade credit (**LQD**), and high-yield credit (**HYG**). Within each regime, asset-specific means and covariances are estimated and used to define six allocation rules that combine classical classical Markowitz method [Markowitz, 1952] with simple heuristics: equal weight, minimum variance, maximum Sharpe ratio, risk parity, a maximum-return corner solution.

1.3 Empirical Design and contribution

The regime-conditioned portfolios are applied in a regime-based monthly backtest from 2007 onward, allowing the strategies to shift risk across equities and bonds as market states evolve. Their performance is benchmarked against a static 60/40 S&P–Treasury portfolio using annualized return, volatility, Sharpe ratio, and drawdown metrics, with particular attention to behavior in stressed regimes.

Our framework enables a direct test of the hypothesis: that a simple, data-driven regime signal derived from basic market and yield-curve information can support an effective dynamic allocation strategy, improving risk-adjusted performance and downside resilience relative to the traditional 60/40 policy mix.

2 Regime Detection

2.1 Method

2.1.1 Feature construction and preprocessing

To detect market regimes we work with monthly data from the McCracken FRED database. From this file we extract the level of the S&P 500 index, the ten–year Treasury yield `GS10` and the three–month Treasury bill yield `TB3MS`. We keep observations from January 1990 onward and convert all numeric fields to floating point values.
The regime signal uses two features. First, we compute the S&P 500 log return

$$r_t^{\mathrm{SPX}} = \log(\mathrm{SPX}_t) - \log(\mathrm{SPX}_{t-1}),$$

which captures monthly equity performance and acts as a direct risk-on or risk-off indicator. Second, we form a term spread as the difference between the long and short Treasury yields,

$$\mathrm{TS}_t = \mathrm{GS10}_t - \mathrm{TB3MS}_t,$$

and then use the monthly change in this spread,

$$\Delta\mathrm{TS}_t = \mathrm{TS}_t - \mathrm{TS}_{t-1},$$

as our second feature. Changes in the term spread summarize shifts in the slope of the yield curve, which are linked to expectations about growth, inflation, and policy moves. We collect these two series in a feature vector $X_t = (r_t^{\mathrm{SPX}}, \Delta\mathrm{TS}_t)$ and drop months with missing values. Before fitting the regime model we standardize both features using `StandardScaler`, so each has mean zero and unit variance. This keeps the hidden Markov model from being driven by differences in scale rather than by genuine co-movements.

2.1.2 Gaussian hidden Markov model

We model the standardized feature sequence $\{X_t\}_{t=1}^{T}$ with a four–state Gaussian hidden Markov model (HMM) implemented in `hmmlearn.GaussianHMM`. The model assumes an unobserved discrete state $S_t \in \{0, 1, 2, 3\}$ that follows a time–homogeneous Markov chain with transition matrix $P = (p_{ij})$, and that the observed feature vector X_t is drawn from a state–specific normal distribution,

$$X_t \mid S_t = k \sim \mathcal{N}(\mu_k, \Sigma_k), \quad k = 0, 1, 2, 3.$$

We set `n_components = 4`, use a full covariance matrix for each state, and allow up to 500 EM iterations with a fixed random seed. Given the standardized input matrix X the model estimates the state means μ_k, covariances Σ_k and the transition probabilities p_{ij} by maximum likelihood.
After fitting the HMM we recover the most likely state sequence $\hat{S}_t$ using the Viterbi algorithm (`hmm.predict(X)` in the code) and store it as `features['state']`. Figure below show the resulting regimes as colored background bands behind the normalized index levels of the S&P 500 and the three bond exchange traded funds (IEF, LQD, HYG). Each contiguous band corresponds to a spell during which the inferred state stays constant. Visually, the bands line up with familiar episodes such as the dot–com boom, the global financial crisis, the post–crisis expansion, the

Covid crash, and the recent inflationary cycle, which supports the economic meaning of the detected regimes.

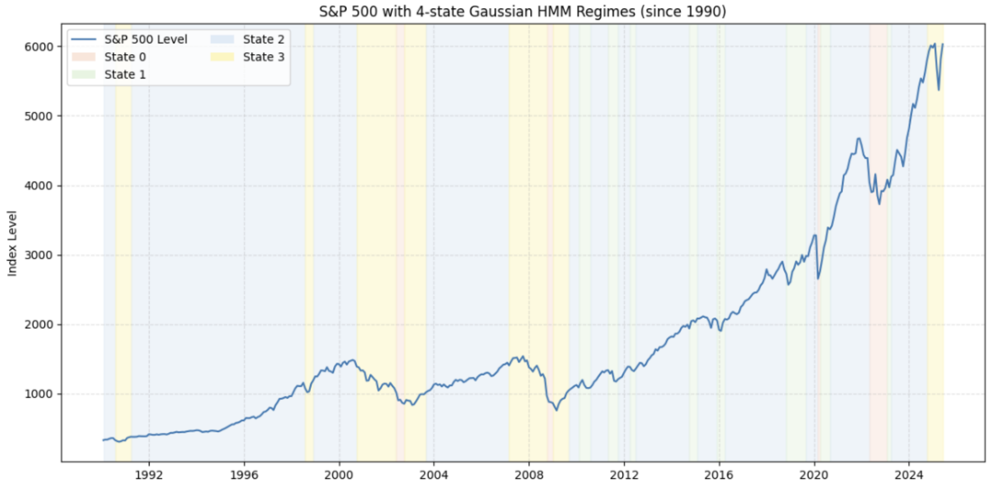

Figure 1: S&P 500 index with four–state Gaussian HMM regimes.

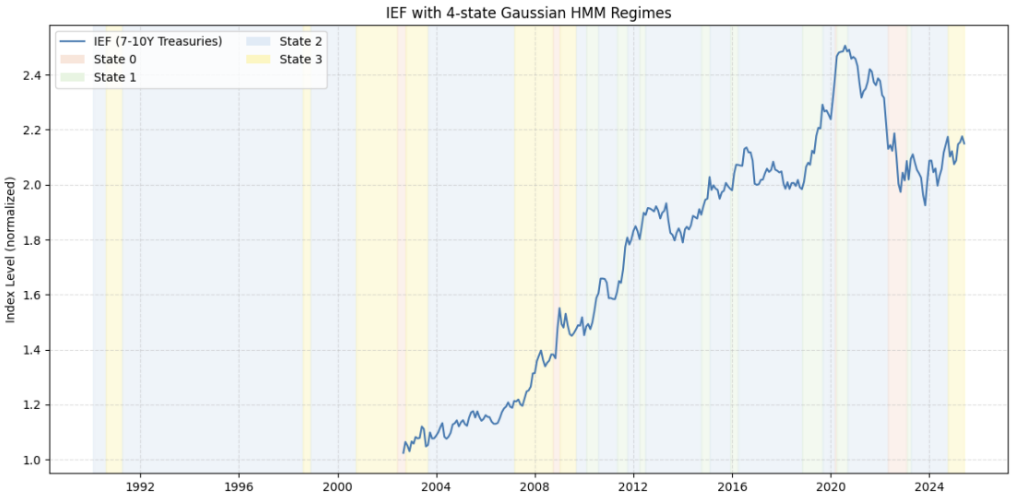

Figure 2: IEF index level with the same four regimes as in Figure 1.

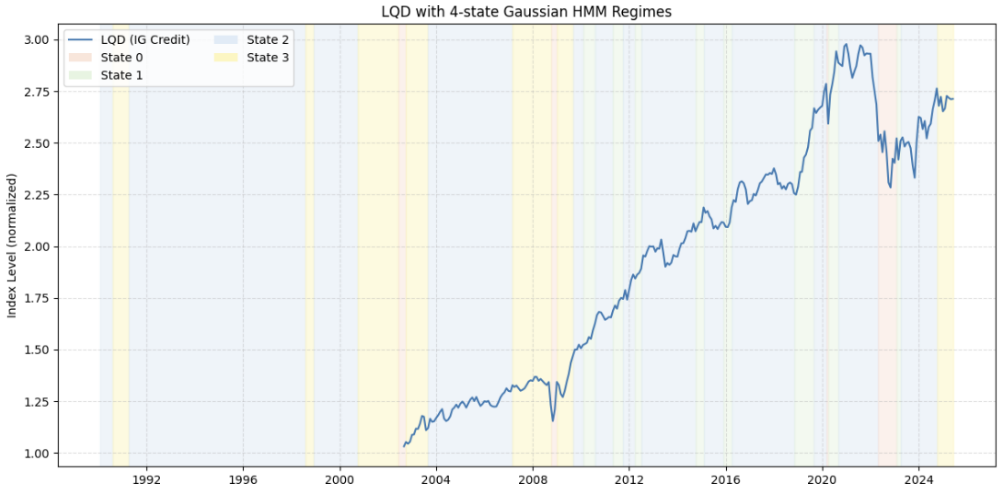

Figure 3: LQD index level with the same four regimes as in Figure 1.

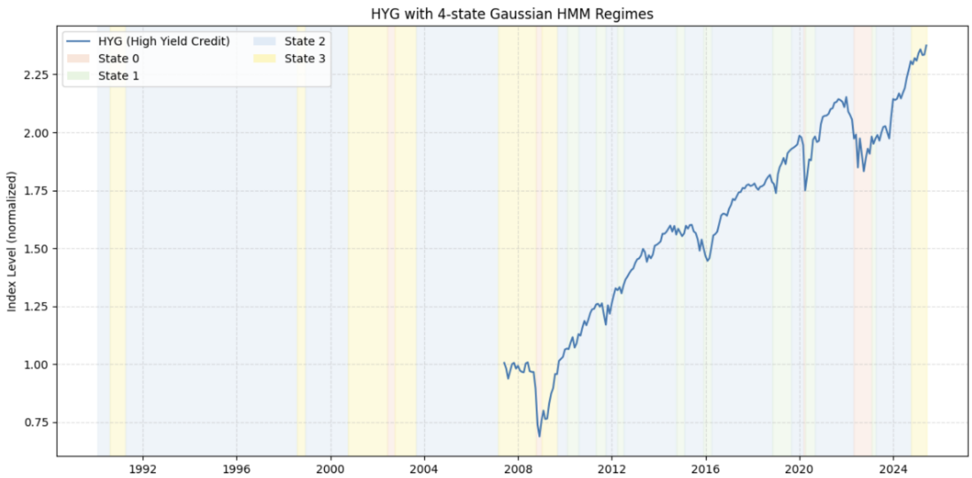

Figure 4: HYG index level with the same four regimes as in Figure 1.

2.2 Regime States Interpretation

To link the latent states to economic conditions we examine how asset returns behave within each state. We first merge the monthly S&P 500 log return and the ETF returns for IEF, LQD and HYG into a panel, align this panel with the inferred state series, and restrict the sample to the period with complete ETF data, starting in June 2007. For each state k we compute the mean and standard deviation of monthly returns for each asset, then annualize these moments and form Sharpe ratios. Table 1 reports the resulting annualized returns, volatilities and Sharpe ratios, together with the number of months assigned to each state.

Table 1: Annualized performance by HMM state.

State	Annualized return SPX	IEF	LQD	HYG	Annualized volatility SPX	IEF	LQD	HYG
0	-0.4553	0.0047	-0.1600	-0.3135	0.3088	0.1241	0.1645	0.2096
1	0.0192	0.1513	0.0968	0.0028	0.1438	0.0598	0.0772	0.0966
2	0.1945	0.0017	0.0429	0.0930	0.0710	0.0575	0.0616	0.0570
3	-0.0022	0.0463	0.0738	0.1399	0.1649	0.0741	0.0877	0.1177

Table 2: Sharpe ratios and number of months by HMM state.

State	Sharpe ratio SPX	IEF	LQD	HYG	Months
0	-1.4745	0.0380	-0.9724	-1.4958	13
1	0.1335	2.5309	1.2526	0.0287	39
2	2.7392	0.0291	0.6965	1.6308	132
3	-0.0133	0.6255	0.8418	1.1889	33

State 0 represents an acute crisis regime. Equity returns are deeply negative—annualizing to nearly minus 45 percent in our sample—and both investment-grade and high-yield credit also post large losses. Treasuries offer only limited offset. This state appears rarely, but it aligns with major stress events when risk assets sell off simultaneously and liquidity evaporates. Examples include the worst months of the 2008 financial crisis and the sudden global market breakdown in March 2020, when forced deleveraging and systemic uncertainty produced indiscriminate selling across equities and credit.

State 1 corresponds to a flight-to-quality or duration-led regime. Equities are roughly flat with a weak Sharpe ratio, yet Treasuries earn about 15 percent annually with extremely low volatility, and investment-grade credit also performs well. High yield, by contrast, barely breaks even. These characteristics point to episodes in which yields fall, duration becomes the main source of return, and investors prefer safer assets without fully abandoning risk. A prominent historical example is 2019, when slowing global manufacturing, a flattening U.S. yield curve, and recurring U.S.–China trade tensions pushed investors into Treasuries. Policy uncertainty and softening growth expectations kept equity markets cautious while duration rallied powerfully.

State 2 reflects a classic risk-on expansion. Equity markets deliver strong, steady gains—the S&P 500 returns nearly 19 percent annually with low volatility and a Sharpe ratio above two—and both investment-grade and high-yield credit perform well. Treasuries earn almost nothing, a pattern consistent with stable or rising yields during periods of firm growth and healthy risk appetite. This type of environment resembles mid-cycle expansions such as 2004–2006 or 2016–2017, when economic conditions were broadly supportive and volatility compressed across asset classes.

State 3 reflects a stressed but not catastrophic environment. Equities are near flat with elevated volatility, while all segments of the bond market earn positive returns and high yield performs the best. The mix suggests slow growth, falling or stable yields, and moderately improving credit conditions. Markets often display this pattern during early-recovery or late-cycle phases, when risk appetite begins to return but equity performance remains uneven. Examples include the post-Eurozone-crisis years of 2012–2013 or the pre-COVID period of 2019 prior to the sharp deterioration, when credit spreads tightened even as equity markets remained sensitive to shifting expectations.

In combination, the feature design and the estimated hidden-state structure allow the model to translate information from stock returns and the yield curve into four economically intuitive regimes. These regimes provide a state variable that naturally supports a dynamic asset-allocation framework, where portfolio weights can adjust in response to changes in the underlying market environment rather than relying on a single static mix.

3 Portfolio Construction Methods

3.1 Investment universe and data frequency

The investment universe consists of four liquid U.S. assets: the S&P 500 total return index, intermediate Treasuries (IEF), investment-grade credit (LQD), and high-yield credit (HYG). All series are converted to **end-of-month** total returns, and portfolios are rebalanced **monthly (regime-based)** in sync with the regime detection process. Optimization rules are applied with a **long-short** strategy, so weights may take negative values when the implied risk-return trade-off favors short exposure to an asset. Portfolio weights are normalized to sum to one at each rebalance date, ensuring full investment while allowing the strategies to express both long and short views across the four asset classes.

3.2 Static 60/40 benchmark

We construct a classic 60/40 stock–bond portfolio to provide a simple and widely understood reference point. The benchmark holds a fixed proportion of capital in the equity index and the Treasury fund IEF. At each rebalancing date the portfolio is realigned to maintain sixty percent in equities and forty percent in Treasuries.
Between rebalancing dates the weights are allowed to drift with relative price movements, which reflects how a buy–and–hold investor naturally becomes more equity heavy in a persistent rally and more bond heavy after an equity drawdown. We consider both monthly and quarterly rebalancing schedules. This benchmark does not use any regime information. It represents the kind of policy portfolio that many investors still follow in practice.

3.3 Regime–based dynamic portfolios

The dynamic strategies use the hidden Markov model state as a compact summary of current market conditions.
Let the state at the beginning of month t be denoted by S_t, and let R_t be the vector of asset returns over that month. Whenever the inferred state changes, the portfolio is allowed to rebalance and update its weights.
To avoid **look–ahead bias** the strategy relies only on returns that precede the decision date and that belong to the same state as the current one.
For each state we collect at most the most recent thirty–six months of historical returns that have been assigned to that state. If there are fewer than twelve such observations, the strategy keeps its previous weights, which prevents unstable estimates in very short samples. Within each state we estimate the average return vector and the covariance matrix of returns. These estimates are then used to define six distinct portfolio rules:

1. **Equal–weight portfolio**. Capital is split evenly across the four assets. With N assets the weight on each asset i is

$$w_i^{\mathrm{EW}} = \frac{1}{N}, \qquad i = 1, \ldots, N.$$

This rule ignores all parameter estimates and treats every asset as equally attractive, which makes it a simple diversification benchmark that is robust to estimation error.

2. **Minimum–variance portfolio**. Here the goal is to minimize predicted portfolio variance while the weights still sum to one. Let Σ be the regime specific covariance matrix and $\mathbf{1}$ a vector of ones. The minimum–variance weights solve

$$\min_{w} \ w^\top \Sigma w \quad \text{subject to} \quad \mathbf{1}^\top w = 1,$$

which in closed form gives

$$w^{\mathrm{MV}} = \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}^\top \Sigma^{-1}\mathbf{1}}.$$

The resulting portfolio leans toward assets that are less volatile and weakly correlated with the rest of the universe.

3. **Maximum–Sharpe portfolio**. Assuming a zero cash rate, we use the regime specific mean vector μ and covariance matrix Σ to construct the tangency portfolio that maximizes the ratio of expected return to volatility. The weights solve

$$\max_{w} \ \frac{w^\top \mu}{\sqrt{w^\top \Sigma w}} \quad \text{subject to} \quad \mathbf{1}^\top w = 1,$$

which implies that the portfolio direction is proportional to $\Sigma^{-1}\mu$. After normalizing we obtain

$$w^{\mathrm{MS}} = \frac{\Sigma^{-1}\mu}{\mathbf{1}^\top \Sigma^{-1}\mu}.$$

Assets with higher estimated returns and favorable covariance with the rest of the portfolio receive more weight.

4. **Risk–parity portfolio**. Instead of balancing expected returns, this strategy balances contributions to total portfolio volatility. Let the portfolio volatility be $\sigma_p(w) = \sqrt{w^\top \Sigma w}$. The risk contribution of asset i is

$$RC_i(w) = \frac{w_i\,(\Sigma w)_i}{\sigma_p(w)}.$$

Risk parity chooses weights such that all contributions are (approximately) equal,

$$RC_1(w) \approx RC_2(w) \approx \cdots \approx RC_N(w) = \frac{\sigma_p(w)}{N}.$$

In practice this leads to higher capital allocations to low–volatility assets and lower allocations to high–volatility assets while still keeping all four exposures in play.

5. **Maximum–return portfolio**. This rule takes an intentionally extreme view. Within the current regime it identifies the asset with the highest estimated mean return

$$j = \arg\max_{i} \mu_i,$$

and allocates the entire portfolio to that asset,

$$w_j^{\mathrm{MR}} = 1, \qquad w_i^{\mathrm{MR}} = 0 \text{ for } i \neq j.$$

The strategy therefore acts as a regime aware “winner takes all” bet and ignores diversification on purpose.

All five strategies are implemented in a fully dynamic way. At each month the current state is observed, the relevant regime–specific weights are applied, and the realized portfolio return is recorded. Between regime changes the portfolios evolve passively with market movements. This design keeps trading focused on times when the hidden Markov model signals that the market environment has shifted in a meaningful way, which is exactly when a regime–based investor would want to reconsider the allocation. The next section evaluates the behavior of these strategies relative to the static 60/40 benchmark.

4 Performance and Analysis of Regime–based Portfolios

Table 3: Annualized performance metrics for benchmark and regime–based portfolios.

Strategy	Ann. return	Ann. vol.	Sharpe	Max drawdown	Win rate
60/40 SPX–IEF (M)	0.0610	0.0880	0.6934	-0.3125	0.6728
60/40 SPX–IEF (Q)	0.0625	0.0880	0.7100	-0.3006	0.6728
SPX Only	0.0793	0.1376	0.5762	-0.5325	0.6636
IEF Only	0.0342	0.0674	0.5077	-0.2319	0.5392
LQD Only	0.0436	0.0792	0.5504	-0.2331	0.5806
HYG Only	0.0539	0.0958	0.5623	-0.3177	0.6498
Regime Equal Weight	0.0526	0.0680	0.7733	-0.2086	0.6452
Regime Min Variance	0.0612	0.0598	1.0233	-0.1675	0.6912
Regime Max Sharpe	0.1095	0.0934	1.1727	-0.2055	0.6498
Regime Risk Parity	0.0541	0.0640	0.8452	-0.1815	0.6498
Regime Max Return	0.0953	0.0956	0.9974	-0.2118	0.6544

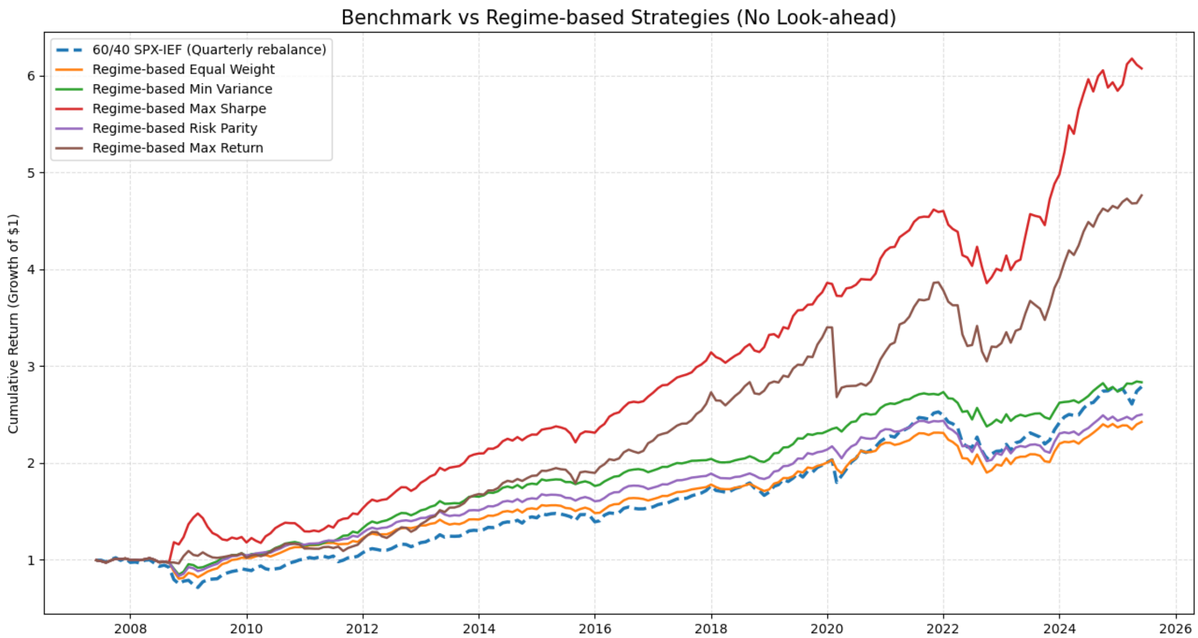

Figure 5: Cumulative wealth for the sixty–forty benchmark and regime-based strategies.

4.1 Performance Analysis

Table 3 reports annualized return, volatility, Sharpe ratio and maximum drawdown for the benchmark portfolios, for the individual assets and for all regime–based strategies. The static sixty–forty portfolio delivers an annualized return of about six percent with volatility close to nine percent and a Sharpe ratio around 0.69. Its maximum drawdown is roughly thirty one percent, which reflects the large equity losses during the global financial crisis and the Covid shock. A pure equity allocation earns a higher average return of almost eight percent but with much higher volatility and a drawdown above fifty percent. In contrast, the Treasury and credit funds have lower returns yet also lower risk, with maximum drawdowns in the low to mid twenty percent range. These single–asset numbers set the stage for the regime–based portfolios, which try to mix the four building blocks in a more flexible way.

The equal–weight regime portfolio has a slightly lower return than the sixty–forty benchmark but also a noticeably lower volatility. Its Sharpe ratio rises to about 0.77 and its maximum drawdown shrinks to around twenty one percent. So simply diversifying across the four assets in a regime–aware way improves risk adjusted performance even without any explicit optimization. The minimum–variance regime portfolio delivers roughly the same return as the sixty–forty mix yet cuts volatility from 8.8 percent to about 6.1 percent. Its Sharpe ratio moves above one and its maximum drawdown improves to about seventeen percent. This strategy shows that using regime specific covariance information to stabilize risk pays off in a very direct way. The risk parity portfolio sit in between. They earn returns just above five percent with volatility a bit above six percent and Sharpe ratios near 0.85. Drawdowns are also smaller than for the benchmark. In other words, these strategies trade some expected return for a smoother ride and better downside protection.
The max return strategy is more aggressive. It earns around 9.5 percent per year with volatility close to that of the benchmark and a Sharpe ratio just below one. Its worst drawdown is still only about twenty one percent. This happens because the strategy tends to concentrate in the asset with the best regime specific mean yet still moves out of harm's way when the regime changes. Finally, the regime based max Sharpe strategy produces the highest raw return at roughly 11.4 percent per year. It does so with much higher volatility near 17 percent and a drawdown of about 43 percent, which brings its Sharpe ratio back in line with the sixty–forty benchmark. Figure 5 visualizes these differences by plotting the cumulative wealth of all regime–based strategies alongside the dashed sixty–forty benchmark. The minimum–variance, risk–parity lines track just above the benchmark with visibly smoother paths, whereas the max Sharpe and max return lines peel away toward the top of the chart but also show sharper dips around major stress episodes. The figure makes the trade off clear: some regime–based strategies use the state information mostly to smooth risk (MinVar, RiskParity), others use it to chase extra return (Max Return and Max Sharpe), and investors can choose where on that spectrum they prefer to sit.

Table 4: Regime specific optimal weights by strategy (ex–post, for interpretation).

Strategy	S&P 500	Treasuries	IG credit	HY credit
Regime 0				
Equal weight	0.250	0.250	0.250	0.250
Minimum variance	0.115	1.016	-0.433	0.303
Maximum Sharpe	0.400	0.252	-0.181	0.530
Risk parity	0.146	0.364	0.274	0.215
Maximum return	0.000	1.000	0.000	0.000
Regime 1				
Equal weight	0.250	0.250	0.250	0.250
Minimum variance	0.077	0.711	-0.159	0.371
Maximum Sharpe	0.032	0.729	0.293	-0.054
Risk parity	0.148	0.356	0.276	0.220
Maximum return	0.000	1.000	0.000	0.000
Regime 2				
Equal weight	0.250	0.250	0.250	0.250
Minimum variance	0.232	0.658	-0.454	0.564
Maximum Sharpe	0.641	-0.019	-0.070	0.448
Risk parity	0.216	0.267	0.249	0.269
Maximum return	1.000	0.000	0.000	0.000
Regime 3				
Equal weight	0.250	0.250	0.250	0.250
Minimum variance	0.048	0.712	-0.108	0.347
Maximum Sharpe	-0.123	0.630	-0.184	0.677
Risk parity	0.154	0.342	0.289	0.215
Maximum return	0.000	0.000	0.000	1.000

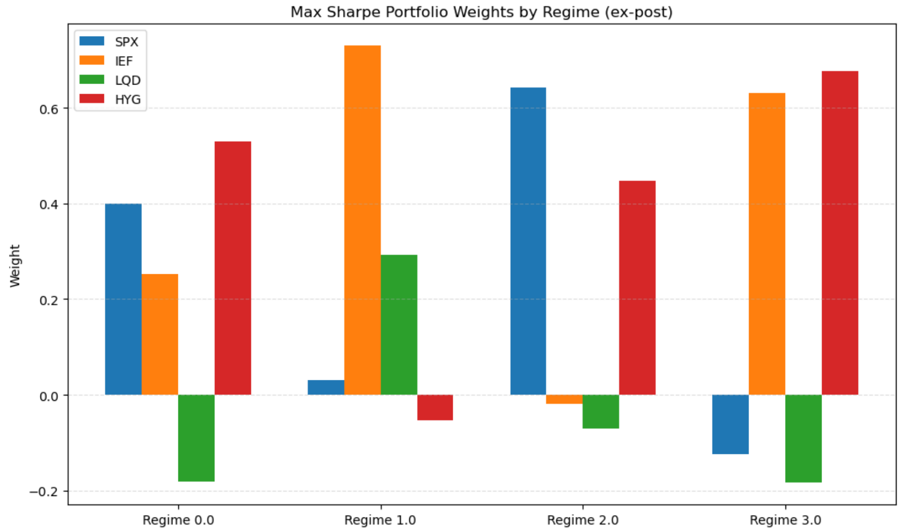

Figure 6: Maximum Sharpe portfolio weights by regime, ex–post.

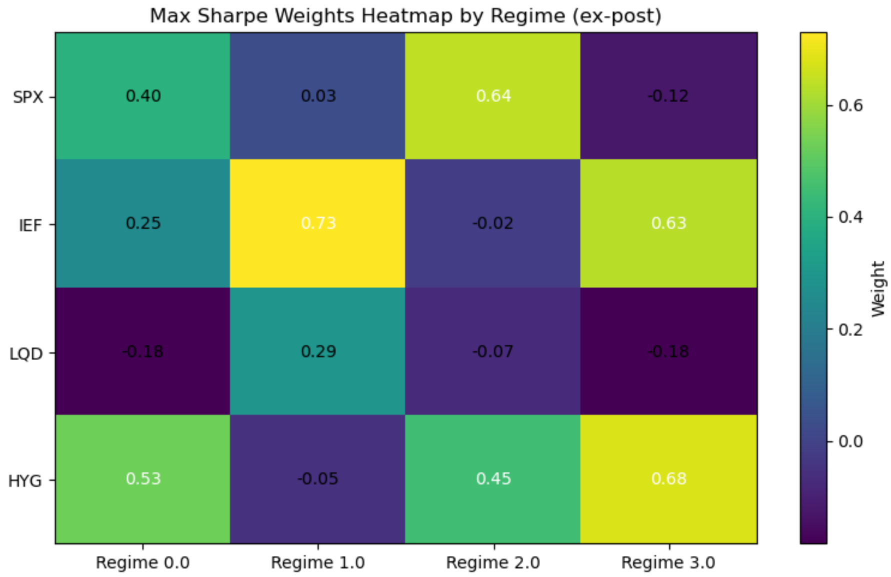

Figure 7: Heat map of maximum Sharpe weights by regime, ex–post.

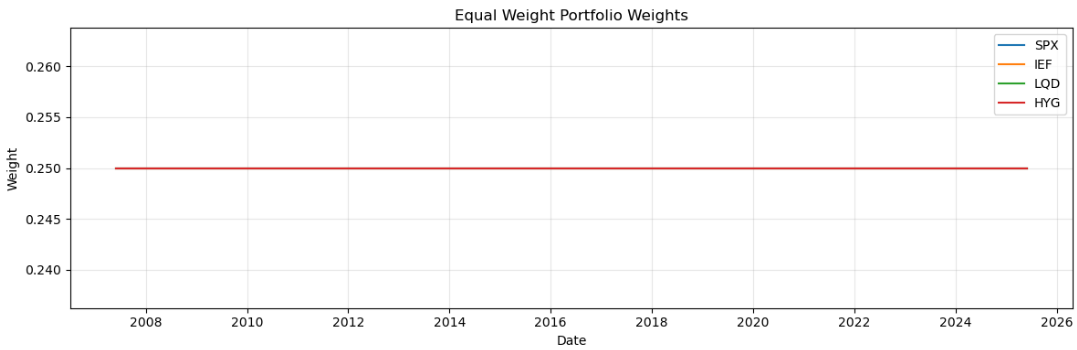

Figure 8: Equal–weight portfolio weights through time.

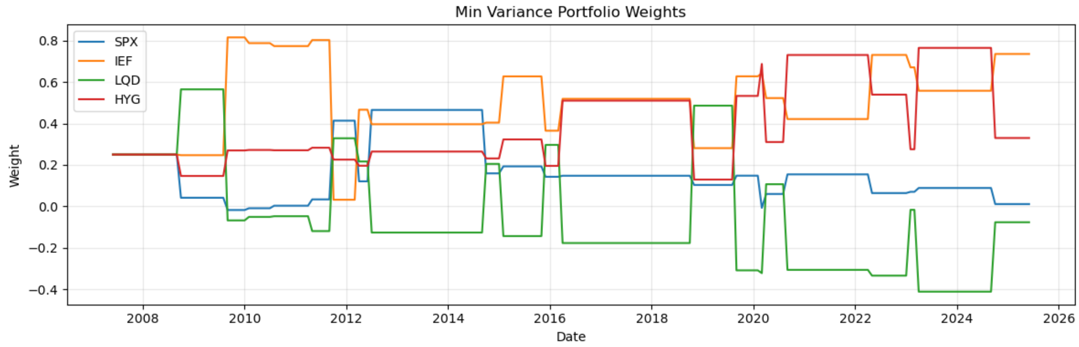

Figure 9: Minimum–variance portfolio weights through time.

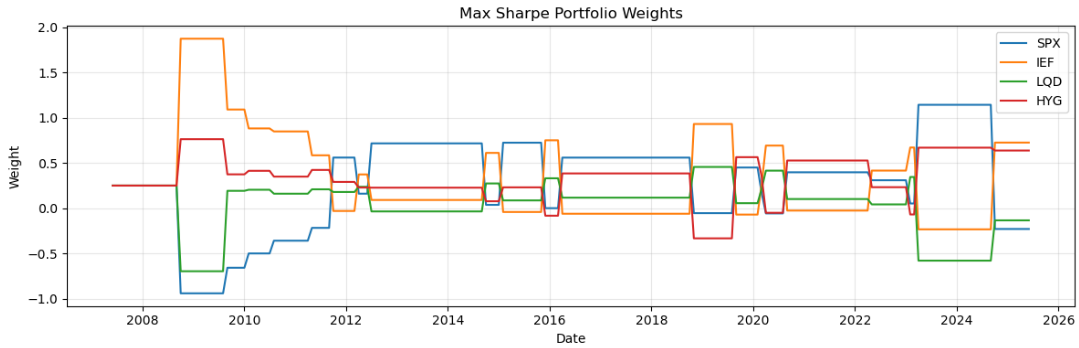

Figure 10: Maximum Sharpe portfolio weights through time.

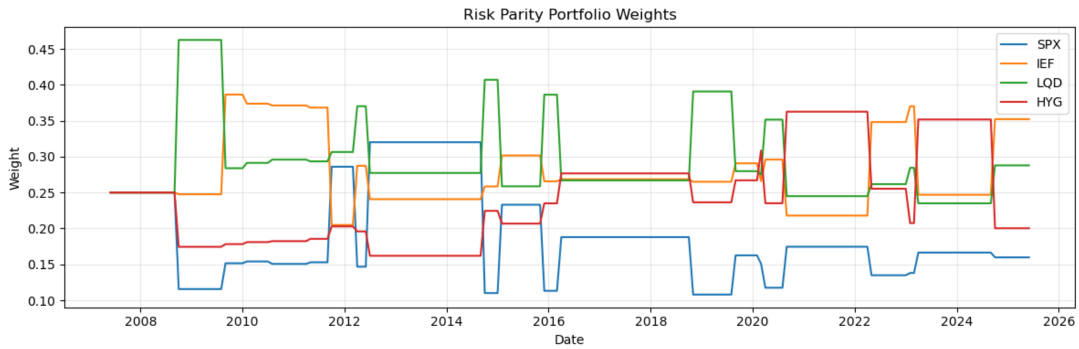

Figure 11: Risk parity portfolio weights through time.

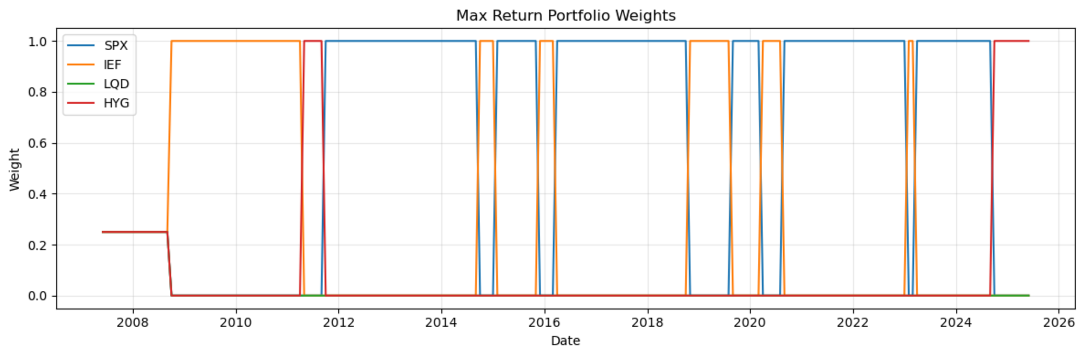

Figure 12: Maximum return portfolio weights through time.

4.2 Regime specific behavior and optimal weights

To understand how the dynamic allocation rules respond to different market conditions, we study the ex-post optimal weights obtained by estimating each portfolio rule separately within each HMM regime using the full sample. These allocations are used only for interpretation. The live back-tests rely on rolling, regime-filtered windows and therefore produce more muted weight changes than the ex-post numbers reported here.

Table 4 summarizes the optimal weights for all strategies across the four regimes. Figures 6 and 7 focus on the maximum-Sharpe portfolio and highlight how its composition shifts across states. In the crisis regime (State 0), the maximum-Sharpe solution tilts to Treasuries and high yield while shorting investment-grade credit. In the flight-to-quality regime (State 1), it concentrates almost entirely in Treasuries. During the expansionary regime (State 2), it loads heavily into equities and high yield, which deliver the highest Sharpe ratios in that state. In the fragile slowdown regime (State 3), the portfolio again over-weights Treasuries and high yield and takes a short position in investment-grade credit.

The time-series of portfolio weights in Figures 8–**??** show how these regime patterns unfold in real time. The equal-weight portfolio remains fixed at one-quarter per asset, providing a stable reference point. The minimum-variance portfolio shifts sharply toward Treasuries during stress periods and occasionally shortens credit, consistent with its low-volatility objective. The maximum-Sharpe portfolio displays the most pronounced variation, with leveraged and sometimes negative exposures that chase the highest expected risk-adjusted return in each regime. Risk-parity portfolio evolve more smoothly: both keep all assets in the portfolio and adjust weights gradually as volatility estimates and state classifications change. The maximum-return rule switches discretely between single-asset allocations and therefore produces a sequence of on/off weight patterns.

These regime-specific weight behaviors help explain the performance results. When the model identifies a risk-on regime, dynamic strategies increase exposure to equities and high yield, improving upside capture. When the model signals crisis or flight-to-quality conditions, allocations shift toward Treasuries, reducing drawdowns and improving Sharpe ratios for the more risk-aware strategies.

5 Conclusion

We investigated whether a regime–based dynamic asset allocation framework improve on the familiar sixty–forty stock–bond mix in a realistic, ETF based setting. To answer it we built a four–state Gaussian hidden Markov model using only two features that investors watch in practice. Monthly S&P 500 returns stand in for broad equity risk, while changes in the Treasury term spread summarize shifts in the yield curve and in macro expectations. The model uncovered four economically intuitive regimes that line up with crisis periods, flight–to–quality episodes, growth phases and more fragile slowdowns. These regimes then became the state variable that drives a family of allocation rules.

On top of this regime signal we layered five portfolio construction methods, ranging from equal weight to minimum variance, maximum Sharpe, risk parity, and maximum return. All strategies invest in the same four liquid U.S. assets and trade at the same monthly frequency. The comparison with the static sixty–forty benchmark shows a clear pattern. Regime–based strategies that use the signal mainly to stabilize risk perform especially well. Minimum variance and risk parity deliver Sharpe ratios above the benchmark and smaller maximum drawdowns, with returns that are similar or only slightly lower. They keep investors in the game while smoothing the ride. More return–seeking rules tell a different but consistent story. The maximum return and maximum Sharpe strategies earn much higher average returns, yet they do so with higher volatility and deeper drawdowns. Regime information helps with time when to take risk, though

it does not eliminate the trade off between growth and pain.
Even a very simple regime model can add value when it is tied to transparent portfolio rules. We did not rely on a long list of predictors, complex machine learning or exotic asset classes. We worked with standard ETFs, monthly data and a basic Gaussian HMM. Yet the combination of states and rules already produced portfolios that look meaningfully different from a static policy mix and that often behave better in terms of risk–adjusted performance and downside protection.
There are certain limitations we need to be aware of. All results are based on U.S. data and on one particular choice of features, so they may not carry over to other markets or other state variables. Transaction costs, financing costs and taxes are ignored. The HMM itself is mis-specified in many ways. Regime means and covariances are estimated with noise, the true return distribution is not Gaussian, and state labels are inferred rather than observed. These caveats matter, yet they do not wipe out the central insight. Even noisy information about where we are in the cycle can support more nuanced allocation decisions.
Future work can push the framework in several directions. One path is to broaden the universe to include global equities, inflation linked bonds, commodities and currencies, so that the regime signal allocates across a richer opportunity set. Another path is to refine the state model by allowing non–Gaussian returns, time varying transition probabilities or alternative machine learning approaches. A third path is to combine regime signals with robust optimization and with explicit trading cost control. All of these extensions share one theme. They try to keep the link between model and portfolio simple enough to understand, yet powerful enough to matter.
In short, the evidence in this paper suggests that regime–based dynamic allocation is a practical step beyond static policy portfolios.
For long–horizon investors who care about both performance and resilience, that balance is exactly what matters.

References

[Abrahamsen and Nakovski, 2021] Abrahamsen, F. and Nakovski, G. (2021). Dynamic asset allocation based on hidden markov model regime sequences. Master's thesis, Norwegian School of Economics.

[Ang and Bekaert, 2002] Ang, A. and Bekaert, G. (2002). International asset allocation with regime shifts. *The Review of Financial Studies*, 15(4):1137–1187.

[Bae et al., 2014] Bae, G. I., Kim, W. C., and Mulvey, J. M. (2014). Dynamic asset allocation for varied financial markets under regime switching framework. *European Journal of Operational Research*, 234(2):450–458.

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

[Markowitz, 1952] Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.

[Nystrup et al., 2018] Nystrup, P., Lindstrøm, E., and Madsen, H. (2018). Dynamic allocation or diversification: A regime-based approach to multiple assets. *The Journal of Portfolio Management*, 44(2):62–76.

[Uysal and Mulvey, 2021] Uysal, A. G. and Mulvey, J. M. (2021). A machine learning approach in regime-switching risk parity portfolios. *The Journal of Financial Data Science*, 3(1):107–127.