10, 11, 12, 13 e 14 de novembro de 2025
ISSN 2177-3866

Multi-Period Optimal Allocation with Regime-Conditional Risk

WANDERCI ALVES BITENCOURT
UNIVERSIDADE FEDERAL DE MINAS GERAIS (UFMG)

ROBERT ALDO IQUIAPAZA
UNIVERSIDADE FEDERAL DE MINAS GERAIS (UFMG)

Agradecimento à orgão de fomento:
We thank FAPEMIG and CNPq.

Multi-Period Optimal Allocation with Regime-Conditional Risk

1. Introduction

Optimal portfolio allocation, traditionally dominated by the Mean-Variance (MV) model of Markowitz (1952, 1959), is widely criticized for its practical limitations (De Jong, 2018; Kolm, Tütüncü & Fabozzi, 2014). These frailties, namely the failure of static models to manage tail risks and abrupt volatility changes, are amplified in emerging markets like Brazil and during periods of financial crisis (Azad & Serletis, 2022; Borio, 2014; Claessens, Kose & Terrones, 2012). This motivated the search for multi-period models, but traditional approaches like Dynamic Programming were often computationally infeasible (Kolm, Tütüncü & Fabozzi, 2014).

A significant advancement in dynamic asset allocation is the integration of Model Predictive Control (MPC) with Hidden Markov Models (HMMs) (Oprisor & Kwon, 2021). MPC provides a tractable framework for adaptive, multi-period optimization (Li, Uysal & Mulvey, 2022; Nystrup et al., 2019), while HMMs excel at identifying latent market regimes and providing regime-conditional parameters (Li, Uysal & Mulvey, 2022; Nystrup, Madsen & Lindström, 2018; Nystrup et al., 2015). This combined approach allows regime information to be dynamically incorporated into the optimal allocation decision.

Despite this progress, the existing literature reveals that most HMM-MPC approaches maintain a fixed risk metric regardless of the market regime. This rigidity prevents the model from fully adapting to the changing nature of investor risk perception, representing a key limitation this study aims to address.

Assuming markets operate under distinct regimes (Ang & Bekaert, 2004; Ang & Timmermann, 2012; Kritzman, Page & Turkington, 2012), a key gap in the literature is the lack of dynamic adjustment in how risk is characterized within an investor's utility function. We address this by proposing a method to modulate risk preferences based on HMM-forecasted volatility regimes. This adaptive approach is fundamental for risk management, especially during periods of high volatility or crisis.

The principle that risk impacts utility non-homogeneously across different market conditions (Jarrow & Li, 2021) provides the theoretical foundation for our approach. By switching the portfolio's risk metric based on the HMM, inferred volatility regime, our model creates a dynamic, state-dependent objective function that better reflects the investor's priorities. This addresses a notable gap in the literature, as few studies have implemented an adaptive switching of risk measures to optimize an investor's long-term expected utility.

This study addresses a key research gap by developing and testing an adaptive, multi-period portfolio optimization strategy. Using a Model Predictive Control (MPC) framework, the model's central innovation is the dynamic switching between *Conditional Value-at-Risk* (CVaR) and *Conditional Drawdown-at-Risk* (CDaR). This switch is conditioned on the market's volatility regime, forecasted by a HMM, thereby operationalizing a state-dependent utility function. Our empirical results from the Brazilian market show that this adaptive strategy significantly outperforms benchmarks that rely on fixed risk metrics.

The choice to alternate between CVaR and CDaR is motivated by the fragility of the stable risk-return premise in volatile markets (DeMiguel, Martín-utrera & Uppal, 2024; Kolm, Tütüncü & Fabozzi, 2014). In low-volatility regimes, CVaR is employed to manage tail risk without the excessive conservatism that could limit gains (Gilbert & Meiklejohn, 2019). Conversely, in high-volatility regimes, CDaR is activated to more effectively limit the risk of accumulated losses (Krokhmal, Uryasev & Zrazhevsky, 2005). This regime-conditional approach aligns the risk objective with the market's temporal structure, aiming to improve risk-adjusted performance by preserving capital in adverse scenarios while capturing opportunities during stable periods.

2. Literature Review

The Mean-Variance (MV) model, introduced by Markowitz (1952), established the foundation of Modern Portfolio Theory (MPT) by formalizing risk as variance (Kolm, Tütüncü & Fabozzi, 2014). Though widely criticized, it remains a primary tool in practice (Tu, 2010).

Key criticisms include parameter estimation errors based on historical data (Michaud, 1989) and, more fundamentally, the absence of a framework for intrinsic uncertainty (De Jong, 2018). These frailties have spurred the development of more robust approaches, including multi-period models and alternative risk metrics (Roman, Mitra & Spagnolo, 2010).

For Kolm, Tütüncü and Fabozzi (2014), the main approaches proposed in the literature to mitigate or circumvent these issues are the Black-Litterman approach (Black & Litterman, 1992), the risk parity approach (Qian, 2005), the equally weighted strategy (DeMiguel, Garlappi & Uppal, 2009), and robust optimization techniques.

Several advanced techniques address the limitations of traditional portfolio models. The Black-Litterman model, for instance, is a shrinkage approach that combines market equilibrium returns with an investor's views (Oprisor & Kwon, 2021). Another strategy, Risk Parity, focuses on risk diversification and avoids the need for return estimation, a critical source of error (Chopra & Ziemba, 2013). Finally, Robust Optimization represents a class of techniques that directly incorporates parameter uncertainty into the problem's formulation (Han, 2020).

The naive, or equally weighted (1/n), allocation strategy is immune to parameter estimation errors and has been shown to frequently outperform sophisticated Mean-Variance models (DeMiguel, Garlappi & Uppal, 2009). According to De Jong (2018), this success stems from the strategy being an optimal way to minimize vulnerability in environments of fundamental uncertainty. On the other hand, a key drawback is that its simplicity leads to an implicit cost: greater exposure to tail risks and extreme events not captured by traditional metrics (Hwang, Xu & In, 2018).

As a method to improve performance under parameter uncertainty, Tu (2010) proposes combining the naive 1/n portfolio with sophisticated optimization strategies. Statistically, this combination acts as a shrinkage estimator that manages the bias-variance trade-off by shrinking the high-variance (though asymptotically unbiased) sophisticated portfolio towards the biased, zero-variance 1/n allocation.

An additional point of debate is that the MV model is inherently single-period, which contrasts with the intertemporal nature of investment decisions in dynamic markets subject to regime changes (Cui et al., 2022; Kolm, Tütüncü & Fabozzi, 2014; Nystrup et al., 2019). The optimal approach for portfolio optimization depends on the investment environment, specifically the level of uncertainty (De Jong, 2018).

The recognition that investment decisions are not restricted to a single point in time, and that investors face transaction costs, intermediate cash flows, and changing market conditions, emphasizes the need to manage risks continuously over the investment horizon and has driven the advancement of multi-period models (Boyd et al., 2017). According to DeMiguel, Martín-utrera and Uppal (2024), the level of market volatility changes over time, which cannot be ignored, much like transaction costs, as both are determinants for achieving outperformance strategies.

While multi-period models based on Dynamic Programming are theoretically robust, their practical application is often infeasible due to severe computational limitations (Kolm, Tütüncü & Fabozzi, 2014). Consequently, the focus has shifted towards more computationally tractable frameworks. Among these, Model Predictive Control (MPC) has proven to be a viable and promising solution for finance (Li, Uysal & Mulvey, 2022; Nystrup et al., 2019; Oprisor & Kwon, 2021). The adoption of such practical approaches, including MPC and the CGL model, is justified by the inherent complexity and technical challenges of solving multi-period optimization problems (Cui et al., 2022).

Two notable approaches to dynamic allocation are the CGL Model and MPC. The CGL Model, from Campani and Garcia (2019), is a specialized method offering an approximate analytical solution that accounts for evolving investor preferences. MPC, on the other hand, is a more general and established control framework, originally from industry, which has been successfully adapted for portfolio optimization in finance.

The application of MPC to multi-period trading was notably explored by Boyd et al. (2017), who used convex optimization to show the framework yields more stable allocations than single-regime models.

MPC is a sequential decision-making approach where, at each time step, an optimization problem over a finite horizon (H periods) is solved, but only the first decision is implemented before the process is repeated with updated information (Boyd et al., 2017). Although potentially suboptimal compared to full-horizon Dynamic Programming, MPC is computationally efficient and allows for the dynamic incorporation of constraints and new forecasts (Nystrup, Madsen & Lindström, 2018). . While advanced applications like drawdown control exist (Nystrup et al., 2019), the literature has typically relied on variance as the risk metric. This overlooks more robust tail-risk measures, such as CVaR or CDaR, which are better suited for crisis regimes.

Furthermore, the realization that markets operate under different regimes has spurred the application of HMM to financial modeling, as introduced by Hamilton (1989). These models constitute a practical and widely implementable approach for capturing the changing nature of return distributions, allowing for the representation of alternating regimes with different statistical characteristics (Cui et al., 2022). Several authors emphasize the importance of considering the dependence of risks and returns on the underlying state of the economy or investor sentiment (Araújo, Camargos & Pinho, 2018; Kolm, Tütüncü & Fabozzi, 2014; Nystrup et al., 2015).

HMMs have become a widely used tool for identifying volatility and return regimes, supporting adaptive decisions based on the transition probabilities between hidden states (Ardia, Bluteau & Rüede, 2019; Bae, Kim & Mulvey, 2014; Peng, Kim & Mittnik, 2022). By incorporating intelligence about market regimes, these models complement multi-period strategies by enabling conditional parameterizations that enrich the dynamic decision-making process.

Recent literature reinforces the need for dynamic, conditional models (DeMiguel, Martín-utrera & Uppal, 2024), spurring growing interest in integrating the predictive power of HMMs with the optimization flexibility of MPC. Seminal works demonstrate the potential of this combination in various forms, such as with mean-variance (Nystrup, Madsen & Lindström, 2018), Risk Parity (Li, Uysal & Mulvey, 2022), or Black-Litterman (Oprisor & Kwon, 2021), grounded in the principle that regime variation affects expected returns (Cui et al., 2022). However, despite these advances, a significant gap remains: these approaches typically employ a fixed risk metric in their objective function. This rigidity limits the model's adaptive capacity, as the characterization of risk itself does not change with the forecasted market state.

This limitation directs focus to the choice of the risk metric itself. Studies have advanced in adopting alternatives to variance, recognizing its inadequacy in crisis or high-volatility scenarios (Alexander & Baptista, 2004). As Ortobelli et al. (2005) point out, there is likely no single ideal risk measure, and the choice can significantly impact performance (Gilbert & Meiklejohn, 2019; Krokhmal, Uryasev & Zrazhevsky, 2005; Lorimer, van Schalkwyk & Szczygielski, 2024).

In this sense, coherent measures focused on the tail of the distribution, such as Conditional Value-at-Risk (CVaR) (Rockafellar, Uryasev et al., 2000) and Conditional Drawdown-at-Risk (CDaR) (Chekhlov, Uryasev & Zabarankin, 2005), have come to stand out. CVaR, although coherent, may have limitations in scenarios with very heavy tails (Gilbert & Meiklejohn, 2019). It is here that CDaR, focused on accumulated losses, emerges as a potentially more conservative and complementary measure, especially in high-volatility regimes where drawdown management is a priority (Harris & Mazibas, 2013).

Still, it is observed that most allocation approaches, even when using these more sophisticated metrics, tend to apply them in a fixed manner over time. Given the evidence that markets exhibit distinct behaviors in different volatility regimes (Ang & Bekaert, 2004; Ang & Timmermann, 2012; Kritzman, Page & Turkington, 2012), a crucial gap remains regarding the evaluation of strategies that dynamically adapt the risk metric itself based on the forecasted market regime. The absence of this adaptation, as suggested by Zhu et al. (2020) in emphasizing the importance of transitioning between scenarios, can lead to a poor assessment of risk.

Additionally, the idea of switching between risk measures finds strong support in state-dependent Utility Theory ($U(x, s)$) (Jarrow & Li, 2021). If the investor's utility is sensitive to the 'state' of the market (volatility regime), then a regime-conditional risk strategy can be seen as a dynamic approximation to maximize this utility, recognizing that different types of losses affect the investor in different ways in different contexts.

Model specification error, where an investor's model fails to capture the true market dynamics (Staden, Dang & Forsyth, 2021), is a critical issue in portfolio choice. A prime example is using a fixed risk metric, which ignores that investor utility is often state-dependent and changes with market regimes. This leads to our central hypothesis: a strategy integrating HMM and MPC that dynamically switches its risk metric (CVaR for low-volatility, CDaR for high-volatility) will outperform equivalent multi-period strategies using fixed risk metrics, both in terms of risk-adjusted returns and expected utility.

3. Methodology

The proposed methodology develops and evaluates a dynamic asset allocation strategy. Its core is a risk metric switch, alternating between CVaR and CDaR, conditioned on the volatility regime identified by a Hidden Markov Model (HMM). The model is trained on a 252-day rolling window to generate out-of-sample forecasts for the next H periods. These forecasts, along with associated return parameters, then feed into a multi-period optimization (MPC) model. Finally, the strategy's performance is benchmarked against fixed-risk and naive strategies using performance metrics and statistical tests.

3. Data

To estimate the Hidden Markov Model (HMM) parameters, we used conditional market volatility (calculated via a GARCH(1,1) model) as the observable variable for the period of January 2019 to April 2025. This choice is supported by literature on market risk and conditions. Due to the 252-day moving window for regime identification, the 2019 data was treated as a warm-up period and excluded from subsequent analyses.

In the allocation stage, the daily returns of a set of ten exchange-traded funds (ETFs) traded on the Brazilian market are used. The analysis period is from January 2020 to April 2025. A more detailed description is presented in Table 1.

As ETF trading in the Brazilian market is relatively recent, data availability for this asset class is still limited. Thus, the choice of the evaluated asset set was based on the criteria of being an asset traded on the B3 and having started trading before January 2020. A second condition is the non-overlapping of the ETFs; that is, if there were two ETFs backed by the same market index, the one with the highest trading volume was chosen.

3. Identification of Volatility Regimes

To identify unobservable market dynamics, we use a HMM with two states (high and low volatility), a choice supported by literature to reduce overfitting risks (Nystrup et al., 2019). The model's observable variable is the conditional volatility of the Ibovespa, estimated via a

Table 1: Data Description

Variable	Description
Exchange-Traded Funds (ETFs):	
BOVA11	Tracks the performance of the Ibovespa/B3 index.
SMAL11	Tracks the performance of the Brazilian *Small Caps* index.
IVVB11	Tracks the performance of the S&P 500 index.
DIVO11	Aims to reflect the performance of the Dividend Index - IDIV/B3.
MATB11	Reflects the performance of the Basic Materials Index - IMAT/B3.
FIND11	Replicates the performance of the Financial Index (IFNC) of B3.
ECOO11	Tracks the performance of the ICO2/B3 index.
GOVE11	Replicates the performance of the B3 Corporate Governance Trade Index (IGCT).
IMAB11	Replicates federal government bonds indexed to the IPCA (IMA-B).
FIXA11	Composed of fixed-rate government bonds, with maturities between 2.5 and 3 years.
Risk-free rate:	
Selic	Short-term interest rate defined by the Central Bank of Brazil.

Note: Daily log-returns of closing prices were used for the analysis. Data for the Ibovespa and ETFs were obtained from Yahoo Finance (https://finance.yahoo.com/), and Selic rate data from the Brazilian Central Bank (BCB) (https://www.bcb.gov.br/). Preprocessing consisted of removing non-trading days to align the series.

GARCH(1,1) process. Using the Expectation-Maximization and Viterbi algorithms on a 252-day moving window, the HMM generates regime-conditioned forecasts for the moments of asset returns (mean and covariance). This forecasting capability is crucial as it allows our asset allocation to adapt directly to the predicted market state.

Thus, the transitions between these two regimes are described by a transition matrix $A = \begin{bmatrix} a_{11} & 1 - a_{11} \\ 1 - a_{22} & a_{22} \end{bmatrix}$, where a_{11} represents the probability of remaining in the low-volatility regime in the next period, given that the previous period was also one of low volatility. Analogously, a_{22} corresponds to the probability of remaining in the high-volatility regime, conditioned on the fact that the previous period was already in that regime.

Assuming that returns follow normal distributions $N(\mu_{11}, \Sigma_{11})$ in the low-volatility regime (LowVol) and $N(\mu_{22}, \Sigma_{22})$ in the high-volatility regime (HighVol), the probability of the market being in the low-volatility regime in period $t + 1$ is calculated based on the probability of being in the LowVol regime in period t and the estimated transition matrix, given by: $\hat{q}_{t+1} = q_t a_{11} + (1 - q_t)(1 - a_{22})$. The forecasted expected return vector for period $t+1$ ($\hat{\mu}_{t+1}$) is a weighted average of the mean returns of the regimes, where the weights are the forecasted probabilities of being in each regime: $\hat{\mu}_{t+1} = \hat{q}_{t+1}\mu_1 + (1 - \hat{q}_{t+1})\mu_2$. The forecast of the return covariance matrix $\hat{\Sigma}$ at time $(t + 1)$ is: $\hat{\Sigma}_{t+1} = (\hat{q}_{t+1}\Sigma_1 + (1 - \hat{q}_{t+1})\Sigma_2) + (\hat{q}_{t+1}(\mu_1 - \hat{\mu}_{t+1})^2 + (1 - \hat{q}_{t+1})(\mu_2 - \hat{\mu}_{t+1})^2)$. The parameters for H periods $(t + 2, ..., t + H)$ can be estimated iteratively, using the estimated transition matrix and the conditional regime forecasts, applied recursively for the following periods.

The GARCH(1,1) model was estimated using the rugarch package and the HMM was fitted with the aid of the depmixS4 package, both in the R language.

3. Multi-period Optimization

Unlike single-period allocation, which determines asset weights by considering the portfolio's performance over time based on a single decision at the initial investment period, otimização Multiperíodo considers that the asset weights for the current period's allocation should be chosen over a planning horizon that extends H periods into the future, from t to $t + H - 1$.

Multi-period optimization (MPO) extends traditional models by planning future allocations to account for transaction costs and anticipated market changes, preventing current decisions from creating future disadvantages. Historically, this approach was hindered by the ”curse of dimensionality” in dynamic programming.

To overcome this, Boyd et al. (2017) developed a tractable MPO model using MPC and convex optimization. In the MPC framework, an optimal plan is created for a future horizon using forecasted parameters, but only the immediate actions for the first period are executed. The problem is then re-solved at the next time step with updated information, a receding horizon technique that simplifies the problem while remaining responsive to market shifts.

Following Li, Uysal and Mulvey (2022) and Oprisor and Kwon (2021), short planning horizons of 5 and 10 periods will be evaluated in order to ensure the quality of the return and risk forecasts, which deteriorates as the prediction horizon increases.

3.3.1 Portfolio with Regime-Conditional Risk

This long-only asset allocation strategy combines HMM forecasting with MPC optimization in a sliding-horizon framework. Crucially, both portfolio rebalancing and the choice of risk measure are conditioned on changes in the market's volatility regime. When a regime change is detected at time t, an H-step optimization problem is solved using the appropriate risk metric for that new regime. In line with the receding horizon approach, only the resulting allocation for the next period (w_{t+1}) is implemented before the entire problem is re-evaluated with updated data on the following day (Boyd et al., 2017; Li, Uysal & Mulvey, 2022; Oprisor & Kwon, 2021).

The generalized optimization formulation by Boyd et al. (2017) can be written in terms of the weights adjusted for an investment horizon H, which is considered relevant for an investment decision to be made at the current date, as:

$$\begin{aligned} \max_{\{w_{t+h|t}\}_{h=1}^{H}} \quad & \sum_{h=1}^{H} \left(\hat{r}_{t+h|t}^{\top} w_{t+h|t} - \lambda \cdot \psi_t^{\text{risk}}(w_{t+h|t}) - \phi \cdot \|w_{t+h|t} - w_{t+h-1|t}\|_1 - \beta \cdot \|w_{t+h|t}\|_2^2 \right) \\ & \text{s.t.} \\ & \quad 1^{\top} w_{t+h|t} = 1, \quad \forall h = 1, \dots, H \\ & \quad w_{t+h|t} \geq 0, \quad \forall h = 1, \dots, H \end{aligned} \tag{1}$$

where $w_{t+h|t} \in \mathrm{R}^N$ represents the asset allocation vector to be decided for each step of the planning horizon $h = 1, \dots, H$, based on the information available at time t. The term $\hat{r}_{t+h|t}$ refers to the expected return vector, while $\psi_t^{\text{risk}}(w_{t+h|t})$ denotes the risk measure conditional on the regime identified at t. The terms $\phi\|w_{t+h|t} - w_{t+h-1|t}\|_1$ e $\beta\|w_{t+h|t}\|_2^2$ penalize, respectively, the transaction cost (changes in weights between periods) and the portfolio concentration, promoting greater diversification. The constraints ensure the full allocation of resources ($\Sigma w = 1$) and the absence of short positions ($w \geq 0$).

We use τ to generically denote each future instant of the decision horizon, i.e., $\tau = t+1, ..., t+H$. Thus, any variable indexed by τ refers to its respective value at step h of the predictive horizon, conditioned on the information available at t.

A central aspect of this model is the switching of the risk function according to the volatility regime estimated via HMM. Thus, the risk function $\psi_\tau^{\text{risk}}(\cdot)$, used at each step of the decision horizon, is defined by:

$$\psi_t^{\text{risk}}(.) = \begin{cases} \hat{\psi}_t^{\text{CVaR}}(\mathbf{w}_\tau), & \text{se o regime identificado for de baixa volatilidade} \\ \hat{\psi}_t^{\text{CDaR}}(\mathbf{w}_\tau), & \text{se o regime identificado for de alta volatilidade} \end{cases} \tag{2}$$

Both functions $\hat{\psi}_t^{\text{CVaR}}$ and $\hat{\psi}_t^{\text{CDaR}}$ are convex functions of w_τ and can be formulated as linear programming problems, which makes the adaptive model computationally feasible for practical applications in moving windows or sequential re-optimization (Chekhlov, Uryasev &

Zabarankin, 2005; Rockafellar, Uryasev et al., 2000).
$CVaR_\alpha$ represents the expectation of losses that exceed VaR_α, and is expressed as:

$$\hat{\psi}t^{\text{CVaR}}(\mathbf{w}_\tau) = \eta + \frac{1}{(1-\alpha)T}\Sigma_{t=1}^{T} \max\left(-R_t(\mathbf{w}_\tau) - \eta, , 0\right) \tag{3}$$

where η is the loss threshold (i.e., the VaR_α) and $R_t(\mathbf{w}_\tau)$ represents the return of the portfolio with weights$\mathbf{w}_\tau$ in period t.
CDaR, in turn, is a measure based on drawdown. Calculated for a confidence level $\alpha \in (0, 1)$, it captures the average of the worst *drawdowns*, i.e., those that exceed the corresponding DaR (*Drawdown-at-Risk*). First, we define:

$$V_{\tau|t}(\mathbf{w}_\tau) = \Sigma_{k=1}^{t} r_k(\mathbf{w}_\tau), \qquad M_{\tau|t}(\mathbf{w}_\tau) = \max_{u \le t} V_u(\mathbf{w}_\tau), \qquad D_{\tau|t}(\mathbf{w}_\tau) = M_{\tau|t}(\mathbf{w}_\tau) - V_{\tau|t}(\mathbf{w}_\tau),$$

where $V_{\tau|t}$ is the cumulative value of the portfolio up to time t, $M_{\tau|t}$ is the historical maximum up to t, and $D_{\tau|t}$ represents the drawdown observed at that point. Thus, the CDaR is given by:

$$\hat{\psi}t^{\text{CDaR}}(\mathbf{w}_\tau) = \eta + \frac{1}{(1-\alpha)T}\Sigma_{t=1}^{T} \max\left(D_t(\mathbf{w}_\tau) - \eta, , 0\right) \tag{4}$$

where η is the *drawdown-at-risk* (DaR) associated with the confidence level α, that is, the smallest value such that $\mathbb{P}(D \le \eta) = \alpha$, or equivalently, $\mathbb{P}(D > \eta) = 1 - \alpha$.
Thus, the aim is to maximize the state-dependent utility function by using an objective function that reflects not only the expected return, but also risk management adapted to the volatility regime, transaction costs, and diversification, ensuring robustness and adherence to the behavior observed in emerging markets.
The state-dependent utility function precludes the use of a fixed risk measure and requires an approach that is state-sensitive and/or distinguishes between different types of risk based on their relationship with the state-dependency in utility, providing the theoretical basis for switching between different risk measures, conditioned on the volatility regime.
Based on recent literature, we selected CVaR for low-instability environments and CDaR for more volatile periods, setting the significance level α to 0.95 for both (Ding, 2023; Harris & Mazibas, 2013; Krokhmal, Uryasev & Zrazhevsky, 2005). This aligns the risk measure with the investor's utility: CVaR addresses the concern for extreme tail losses in stable markets, while CDaR focuses on mitigating the discomfort of continuous, accumulated drawdowns during turbulent periods, which can disproportionately impact an investor's sense of wealth preservation (Daniel, 2025).
This alternation between CVaR and CDaR is a sophisticated response to the limitations of variance as a risk measure, which ignores the tails of financial return distributions. The convex optimization problem described in Equation 1, which incorporates these risk metrics, was implemented and solved using the CVXR package in the R language, with the ECOS (Embedded Cone Solver) solver.
The proposed multi-period allocation approach with regime-conditional risk is compared to its unconditional-risk counterparts for both multi-period and single-regime allocations, as well as a naive, equally weighted allocation ($1/n$).

3. Definition of Model Parameters

To mitigate the estimation risks inherent in portfolio optimization, which often lead to poor out-of-sample performance, we adopt a data-driven calibration approach, considered essential (Han, 2020). Data-driven calibration seeks to determine parameter values that mitigate the impact of estimation errors in the observed data (Kircher & Rösch, 2021).

In this work, the optimization hyperparameters, risk aversion (λ), concentration penalty (β), and transaction cost (ϕ), were systematically tuned using the irace algorithm. Irace finds the optimal parameter configuration by maximizing an objective function over a set of simulations. As our objective function, we use the net Certainty Equivalent Return (CER), a measure of investor utility that evaluates out-of-sample performance. The net CER is defined as:

$$r_{p,t}^{\text{gross}} = \Sigma_{i=1}^{n} w_{i,t} \cdot r_{i,t} \tag{5}$$

$$r_{p,t}^{\text{net}} = r_{p,t}^{gross} \cdot (1 - \phi \cdot \|w_t - w_{t-1}\|_1), \tag{6}$$

and the CER as:

$$CER = r_{p,t}^{\text{net}} - \frac{\lambda}{2} \cdot \sigma^2(r_{p,t}^{\text{net}}), \tag{7}$$

where w_t are the portfolio weights, $r_{p,t}^{\text{gross}}$ are the gross portfolio returns, $r_{p,t}^{\text{net}}$ are the net portfolio returns, and $\sigma^2(r_{p,t}^{\text{net}})$ is the variance of the net returns. This automated and reproducible process (ensured by a fixed random seed) determines the portfolio weights for each date, dynamically applying the risk metric (CVaR or CDaR) according to the volatility regime forecasted by the HMM model.

The hyperparameter search space was defined in light of the empirical results and suggestions found in recent works (Campani, Garcia & Lewin, 2021; Daniel, 2025; Nystrup et al., 2019; Oprisor & Kwon, 2021). The internal parameters of Irace followed the standard recommendations from the literature (for details, see Lopez-Ibanez et al. (2016)).

3. Evaluation Procedures and Performance Metrics

We compared the performance of all strategies and benchmarks using a comprehensive set of metrics for return, absolute/relative risk, and risk-adjusted efficiency. Statistical significance of the differences was then assessed with a non-parametric test.

For absolute risk, we used Annualized Volatility, calculated as the daily standard deviation of returns scaled by $\sqrt{252}$, and Maximum Drawdown (MDD). To measure rebalancing activity and implicit transaction costs, we calculated portfolio Turnover using the following expression:

$$\text{Turnover} = 1/(T-1) \cdot \Sigma_{t=1}^{T-1} \Sigma_{i=1}^{N} |w_{i,t} - w_{i,t-1}|, \tag{8}$$

where $w_{i,t}$ represents the weight of asset i at time t.

Relative risk was measured using the Tracking Error (TE), used to evaluate the active deviation of a strategy relative to a benchmark. The TE was calculated as the annualized standard deviation of a difference between the daily portfolio returns (R_t^{estrat}) and the benchmark returns (R_t^{bench}):

$$TE = \sqrt{\frac{1}{n-1} \Sigma_{t=1}^{n} (R_t^{\text{estrat}} - R_t^{\text{bench}})^2} \times \sqrt{252} \tag{9}$$

To measure risk-adjusted return, the Sharpe, Sortino, and Omega ratios were employed. The Sharpe Ratio was calculated based on the mean and standard deviation of the returns over the entire analysis period, considering the correction suggested by Israelsen (2005).

Due to the non-normal character of the return distributions, we conducted statistical significance tests using the Aligned Rank Transform (ART), a non-parametric method that enables the use of ANOVA on ranked data for factorial designs (Wobbrock et al., 2011).

In this study, ART was applied to the daily Sharpe Ratio time series from a 252-day moving window. Using the allocation strategy as the main factor, we employed the Tukey test for multiple post-hoc comparisons, implemented in R with the `ARTool` package.

4. Empirical Results

This section discusses the empirical results for our dynamic allocation strategy, which uses an HMM-MPC framework to switch between CVaR and CDaR based on the prevailing volatility regime. The model was evaluated using historical data of ETFs traded on the B3 over a period encompassing different economic cycles and market conditions.

Figure 1. Cumulative return paths for the evaluated strategies. The HMM-MPC-Adaptive strategy is the main proposal, which features a dynamic risk metric that switches based on the HMM-inferred volatility regime.

For comparison, we evaluated several benchmarks. First, two multi-period strategies that use HMM regime forecasts but a fixed risk metric: the HMM-MPC-CVaR and the HMM-MPC-CDaR. We also tested two single-period models that use HMM forecasts but lack the multi-period optimization component: the HMM-Single-CVaR and the HMM-Single-CDaR. Finally, the naive Equal Weight (EW) portfolio (1/n allocation) was included as a standard literature benchmark.

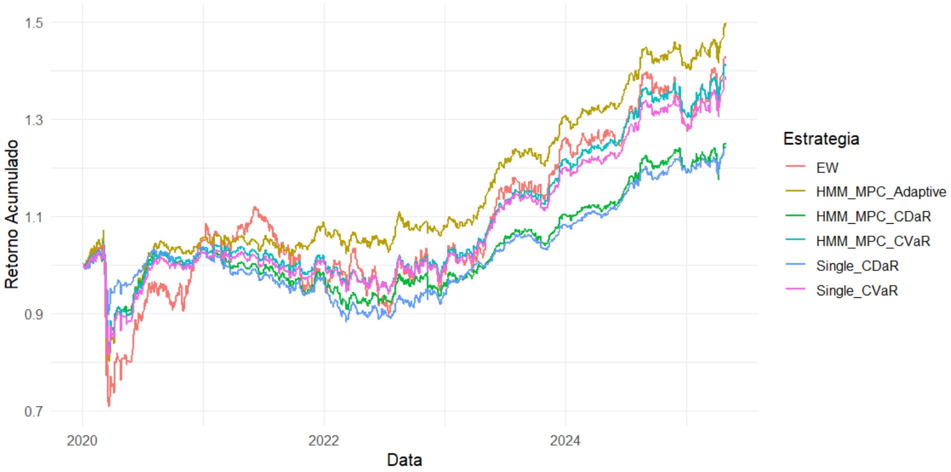

Figure 1: Cumulative return of the strategies for the period from January 2020 to April 2025 and a planning horizon of 5.

All strategies, with the exception of the naive portfolio (EW), use the same set of hyperparameters: $\lambda = 3$ (risk aversion), $\phi = 0.2\%$ (proportional transaction cost), and $\beta = 0.1$ (concentration penalty). These values were selected in a *data-driven* manner by the irace algorithm, ensuring that any performance difference arises solely from the approach proposed by the evaluated investment strategy.

As observed in Figure 1, the HMM-MPC-Adaptive strategy shows a superior performance, in terms of cumulative return, compared to the other strategies throughout the sample period. This graphical evidence is confirmed by the metrics summarized in Table 2, which synthesize the performance of the different approaches.

Table 2 compares planning horizons of $H = 5$ and $H = 10$. We chose $H = 5$ as the default configuration, as a robustness test with $H = 10$ did not significantly alter the conclusions. This decision is supported by literature indicating that shorter horizons reduce prediction error (Li, Uysal & Mulvey, 2022) and have little impact on frequently rebalanced strategies (Campani, Garcia & Lewin, 2021). Furthermore, the $H = 5$ model demonstrated superior performance in terms of Certainty Equivalent Return (CER), despite a slightly lower gross return.

As indicated in Table 2, the HMM-MPC-Adaptive strategy shows superiority across several performance metrics. For the period from January 2020 to April 2025, the multi-period strategy

Table 2: Performance metrics of the strategies from January 2020 to April 2025.

Strategy	Cum.Ret.	Ann.Ret.	Sharpe	Sortino	Omega	MDD	Vol.	Turn.	CER
Multi-period with H = 5:									
HMM-MPC-Adaptive	48.63	8.84	**0.022**	**0.071**	**1.20**	20.58	10.74	1.17	**7.10**
HMM-MPC-CDaR	37.18	6.99	0.013	0.063	1.18	18.78	9.68	1.12	5.49
HMM-MPC-CVaR	41.47	7.70	0.016	0.064	1.17	19.47	10.43	0.58	6.34
Multi-period with H = 10:									
HMM-MPC-Adaptive	**49.79**	**9.02**	**0.022**	0.065	**1.20**	25.14	11.92	1.30	6.96
HMM-MPC-CDaR	25.09	4.90	0.001	0.041	1.12	21.37	10.62	1.05	3.25
HMM-MPC-CVaR	41.34	7.68	0.016	0.060	1.17	21.34	10.98	0.59	6.21
Single-regime:									
HMM-Single-CDaR	24.39	4.78	-0.002	0.052	1.13	**15.30**	**8.13**	1.18	3.51
HMM-Single-CVaR	38.70	7.24	0.014	0.057	1.16	21.89	11.00	0.61	5.75
Naive benchmark:									
EW	42.91	7.93	0.014	0.043	1.10	32.34	18.04	**0.00**	4.85

Note: Bold values indicate the best result in the column (maximum or minimum, as applicable). All strategies only allow for long positions. Cum.Ret. and Ann.Ret. are abbreviations for cumulative return and annualized return, respectively. MDD is the abbreviation for Maximum Drawdown, Vol. for volatility, and Turn. for turnover. The values for Cum.Ret., Ann.Ret., MDD, Vol., and CER are presented in percentage terms. Sharpe refers to the adjusted version proposed by Israelsen (2005).

with regime-conditional risk delivered a cumulative return of 49%, outperforming the $1/n$ portfolio (43%).

The multi-period strategies with a regime-unconditional risk metric also exhibited higher cumulative returns than their single-regime counterparts; however, the relative gain is more modest when comparing the difference with the conditional-risk approach. These results corroborate the literature that highlights the relevance of strategies capable of adapting to the structural changes in financial markets (Li, Uysal & Mulvey, 2022; Oprisor & Kwon, 2021).

Switching between CVaR in normal regimes and CDaR in turbulent ones proved effective, balancing the capture of gains in low-volatility periods with capital protection during high-volatility drawdowns. These results reinforce the thesis that there is no universally optimal risk metric (Ortobelli et al., 2005) and that adapting the risk measure to the prevailing market regime leads to superior risk-return outcomes and greater investor utility (Lorimer, van Schalkwyk & Szczygielski, 2024).

In contrast, the multi-period strategies using a fixed (unconditional) risk metric failed to outperform the naive portfolio. This underperformance can be attributed to transaction costs; the benefits gained from using sophisticated forecasting and risk models were insufficient to overcome the costs associated with the strategies' higher turnover.

In summary, the HMM-MPC models with unconditional risk outperform single-regime portfolios, but not the adaptive model or the naive one, emphasizing that the success of the strategy depends on both the multi-period horizon and the adopted risk metric (Gilbert & Meiklejohn, 2019). The adaptive advantage also validates the relevance of tail risk metrics in markets with non-normal returns (Harris & Mazibas, 2013).

Regarding risk-adjusted return, the HMM-MPC-Adaptive maintains its advantage. The $1/n$ portfolio, although robust to estimation errors (Tu, 2010), exhibits a Sharpe Ratio (0.014) lower than the HMM-MPC-Adaptive and HMM-MPC-CVaR strategies. The Omega Ratio, although similar among the multi-period strategies, also confirms the superiority of the HMM-MPC-Adaptive model, indicating a gain-to-loss ratio of 20% ($\Omega = 1.20$).

The CER criterion, used in this study both for parameter calibration and for out-of-sample performance evaluation, represents the certain return that an investor with a quadratic utility function ($U(R) = E[R] - \frac{\lambda}{2}Var(R)$) would consider equivalent to an investment with a risky return. The use of CER as an out-of-sample evaluation criterion ensures theoretical consistency, as it directly reflects the optimization objective of Equation 1, and according to this criterion, the HMM-MPC-Adaptive strategy obtained the highest CER (7.10%) among all evaluated strategies. It is observed that the multi-period approach, by employing regime-conditional risk, adds 0.76 p.p. of CER over the multi-period approach with regime-unconditional CVaR risk, 1.61 p.p. of CER over the multi-period approach with regime-unconditional CDaR risk, and 2.25 p.p. over the equally weighted portfolio.

The superior Certainty Equivalent Return (CER) of the adaptive strategy empirically confirms its effectiveness in maximizing the investor's risk-adjusted utility. A higher CER is a key indicator of a strategy's ability to offer a more favorable balance between potential gains and losses.

In terms of absolute risk, the HMM-MPC-Adaptive strategy's Maximum Drawdown and volatility are comparable to other multi-period models and significantly lower than the 1/n strategy (10.74% vs. 18.04%). Achieving the highest CER under these conditions demonstrates that adapting the risk metric effectively improves the risk-return trade-off without sacrificing capital protection.

Turnover, a crucial metric reflecting management intensity and transaction costs, was lowest for the CVaR-only strategy. The higher turnover of the adaptive and CDaR-only models stems from CDaR's conservatism and the dynamic weight changes inherent in the adaptive approach. While all HMM-based models rebalance on regime changes, only the adaptive strategy also switches the risk metric, contributing to its active management profile.

The TE shows how far the path of an active strategy deviates from its benchmark. According to Table 3, the HMM-MPC-Adaptive maintains a TE of 10.37% against the $1/n$ portfolio, indicating that the regime-conditional rebalancing mechanism delivers superior returns without producing excessive active volatility. This consistency reinforces the efficiency of the approach, which already prospectively internalizes transaction costs in the planning horizon, and confirms the thesis that switching the risk measure conditional on the regime improves the risk-return *trade-off*.

The CDaR-guided portfolios exhibit the largest Tracking Error (TE), reaching up to 13.01% relative to the 1/n. This deviation stems from two combined factors: CDaR's path-dependent nature, which forces deep risk cuts after prolonged drawdowns, and the abrupt reallocations following HMM regime-change signals. A high TE, however, does not imply lower quality. For instance, the HMM-MPC-CDaR model maintains a lower MDD and volatility than the adaptive approach. This illustrates that a more defensive, drawdown-focused strategy can improve capital preservation at the cost of greater divergence from a benchmark, a finding consistent with literature that distinguishes drawdown-based metrics from those focused only on punctual tail losses (Chekhlov, Uryasev & Zabarankin, 2005; Krokhmal, Uryasev & Zrazhevsky, 2005).

Table 3: Tracking Error between pairs of strategies.

	Benchmark				
Strategy	HMM-MPC CDaR	HMM-MPC CVaR	HMM-Single CDaR	HMM-Single CVaR	Equal Weight
HMM-MPC-Adaptive	3.23	3.06	5.80	3.20	10.37
HMM-MPC-CDaR	–	3.47	4.58	4.15	11.84
HMM-MPC-CVaR	–	–	4.67	1.96	9.84
HMM-Single-CDaR	–	–	–	5.39	13.01
HMM-Single-CVaR	–	–	–	–	9.28

Note: The following symmetric matrix was calculated from daily portfolio returns. Only the upper triangle is displayed.

An Aligned Rank Transform (ART) with a fixed-effects ANOVA was applied to the daily Sharpe Ratios to test for statistical significance. The global test was highly significant ($F = 458.43$, $p < 2.2 \cdot 10^{-16}$), allowing us to reject the null hypothesis of equal performance. This result confirms that at least one strategy's risk-adjusted return distribution differs from the others, justifying post-hoc comparisons.

The results of the Tukey's post-hoc test, with corrections for multiple comparisons, demonstrate the robust superiority of the HMM-MPC-Adaptive strategy. It outperformed all other strategies with statistical significance at the 1% level for all pairwise comparisons. This provides compelling empirical evidence that a multi-period, regime-conditional risk approach generates superior risk-adjusted returns, thereby validating this study's central hypothesis.

To deepen the understanding of the underlying mechanisms behind the superior performance of the HMM-MPC-Adaptive strategy, an analysis of the portfolio's behavior conditioned on the volatility regime identified by the HMM was performed, and, consequently, on the active risk metric (CVaR in low-volatility regimes and CDaR in high-volatility regimes). Table 4 summarizes the descriptive statistics of the daily portfolio returns under each of these conditions.

Table 4: Statistics of the HMM-MPC-Adaptive portfolio, conditioned on the active risk metric.

Risk	Obs.	Mean	Median	SD	Min.	Max.	Perc_5	Perc_95	Skew.	Kurt.	MDD
CVaR	858	7e-04	1e-03	0.008	-0.034	0.029	-0.012	0.013	-0.159	0.937	0.123
CDaR	369	-4e-04	2e-04	0.017	-0.121	0.077	-0.022	0.022	-1.645	11.822	0.321

Note: The statistics are calculated on daily returns. Obs. represents the number of days each risk metric was active. SD, Min, Max, Skew., and Kurt. are the standard deviation, minimum return, maximum return, skewness, and kurtosis. P5 and P95 are the 5th and 95th percentiles, respectively. MDD is the Maximum Drawdown (the largest peak-to-trough percentage loss) observed within a continuous period where the respective risk metric was active.

It is observed that the strategy operated under the CVaR logic most of the time, approximately 70% of the backtesting period, indicating a predominance of low-volatility regimes identified by the HMM.

The daily return dynamics (Table 4) corroborate the logic of the switching mechanism. During the low-volatility periods when CVaR was active, the portfolio successfully captured opportunities, delivering a positive average daily return of 0.067% (median 0.098%) with a volatility of 0.805%. The return distribution was nearly symmetric (skewness -0.159) with moderately heavy tails (kurtosis 0.94). These results, including a worst intra-regime drawdown of 12.3%, confirm that CVaR allowed for exposure to gains without excessive conservatism, fulfilling its intended function in stable regimes.

In contrast, during the high-volatility periods when CDaR was active, the market environment was more challenging, with an average daily return of -0.036% and daily volatility jumping to 1.75%. The return distribution showed strong negative skewness (-1.65) and high kurtosis (11.8), with extreme daily movements from -12.2% to +7.74%, a profile typical of markets in crisis. However, the strategy still achieved a positive median return (0.025%) and captured significant gains on recovery days. This highlights a key feature of the CDaR-guided approach: it provides protection without becoming so restrictive that it misses crucial market rebounds.

The main justification for using CDaR in high-volatility regimes, however, lies in its ability to limit prolonged drawdowns. The analysis of the worst continuous intra-regime drawdown (Table 4) reveals that during the phases when CDaR was active, the portfolio faced a maximum drawdown of 32.1%. Notably, this value coincides with the largest overall drawdown suffered by the HMM-MPC-Adaptive strategy, which occurred during the peak of the COVID-19 crisis in March 2020, a period when CDaR was active, as predicted by the HMM. Although daily volatility and punctual losses were high, reflecting the severity of the market environment, the effectiveness of CDaR in managing the depth of accumulated loss becomes evident when compared to benchmarks. During this same critical period of March 2020, the Ibovespa suffered a drawdown of 50.24%.

The conditional analysis in Table 4 provides empirical evidence for the effectiveness of the HMM-MPC-Adaptive strategy's design. During low-volatility regimes, the active CVaR metric successfully captured opportunities, generating positive average returns and driving the overall cumulative gain. Conversely, during high-volatility regimes, the active CDaR metric proved effective for portfolio protection. By focusing on capital preservation during these turbulent phases, the control over accumulated losses was crucial to achieving the strategy's superior overall risk-adjusted performance.

The conditional analysis reveals why the HMM-MPC-Adaptive strategy succeeds where fixed-risk benchmarks fail. Its design of switching between CVaR in low-volatility periods and CDaR in high-volatility ones provides a crucial adaptive balance. This allows the portfolio to capture opportunities in stable markets while ensuring robust protection against accumulated losses during crises,a balance that is lost when either risk metric is used unconditionally.

Beyond practical gains, the proposed model addresses a structural deficiency in traditional allocation: model misspecification error. By integrating a multi-period MPC framework with HMM-guided risk switching, our approach operationalizes a state-dependent utility function. As discussed by Staden, Dang and Forsyth (2021), this tackles a different challenge than parameter error; our model addresses the error of assuming an investor's utility function like static. By switching between CVaR and CDaR according to the market regime, the model adapts its risk specification to be most pertinent to the current market state, thus mitigating losses from a rigid and inadequate risk framework.

Although our model does not explicitly predict what the 'true model' of returns will be at a given moment, the adaptation of the risk metric reflects the investor's perception of the prevailing risk profile, which is shaped by different tail characteristics of the return distribution, characteristics that would be represented by distinct underlying models (e.g., a more symmetric distribution in low volatility versus heavy tails in high volatility).

This approach also dialogues with recent literature in finance, which recognizes the relevance of adaptive strategies in the face of the heterogeneity of volatility regimes (DeMiguel, Martín-utrera & Uppal, 2024). As the results show, adapting the risk metric not only improves the adjusted return but also contributes to the stability of the allocation in the face of structural changes, enhancing the investor's net expected utility, evidence captured by the superiority of the Certainty Equivalent Return.

Therefore, by allowing the objective function to dynamically mold itself to the prevailing risk regime, the HMM-MPC-Adaptive model offers a practical and effective solution to the challenge of asset allocation in markets with intertemporal volatility. It represents a relevant theoretical and empirical contribution, with direct implications for the practice of portfolio management, especially in emerging economies.

5. Conclusion

This study proposes a dynamic asset allocation model using MPC and HMM, featuring a novel regime-conditional risk metric. The model's objective function switches between CVaR and CDaR based on the HMM-identified volatility regime, filling a literature gap where even multi-period models use a fixed risk measure and thus fail to fully adapt to market changes.

The main theoretical contribution is the operationalization of a state-dependent utility function via a practical optimization approach. Instead of modeling the utility function directly, our model dynamically switches the risk metric in its objective function based on the HMM-inferred regime. The switch between CVaR in low-volatility periods and CDaR in high-volatility ones acts as a proxy for an investor's changing sensitivity to different types of losses (punctual vs. accumulated). This dynamic adaptation is interpreted as a mechanism against the model misspecification error of assuming a static utility function, ultimately aiming to enhance the expected utility of wealth.

The empirical tests used a representative set of ETFs listed on the B3 between January 2020 and April 2025, a period capturing the COVID-19 crisis and multiple volatility transitions in Brazil. The HMM-MPC-Adaptive strategy systematically outperformed all considered benchmarks, including single-regime portfolios and multi-period variants with fixed risk. This superiority was confirmed by the ART/ANOVA test, which showed statistically significant differences in risk-adjusted returns.

The adaptive strategy's superior Certainty Equivalent Return (7.10%) suggests it offers a higher expected utility, even after accounting for transaction costs. Additionally, the strategy maintained a level of absolute risk (Maximum Drawdown and volatility) comparable to other multi-period models and lower than the naive 1/n strategy, reinforcing that adapting the risk measure improves the risk-return balance without sacrificing protection. The portfolio's turnover, although higher than some benchmarks, reflects the active and adaptive nature of the management.

Our theoretical contributions are twofold. First, the effectiveness of switching between CVaR and CDaR validates the concept of regime-dependent risk utility, confirming no single metric is universally optimal (Gilbert & Meiklejohn, 2019; Ortobelli et al., 2005). Second, the success of the strategy with a short planning horizon (H=5) highlights that the frequency of rebalancing and the quality of short-term forecasts are more critical than the horizon's length, a finding consistent with both Campani, Garcia and Lewin (2021) and DeMiguel, Martín-utrera and Uppal (2024).

From a practical standpoint, this study validates adaptive strategies that switch risk metrics based on market regimes, making the proposed HMM-MPC framework a promising alternative for portfolio management, especially in volatile emerging markets. By simultaneously controlling drawdowns and enhancing risk-adjusted returns, this adaptive architecture represents a significant improvement over prior models that used a fixed risk metric.

Despite its contributions, this study has limitations, primarily its restriction to ETFs on the Brazilian market, which limits the generalizability of the results. Future research could extend this work by applying the methodology to other markets, testing alternative risk metrics like Entropic Value at Risk, and exploring more sophisticated transaction cost models.

References

Alexander, G. J., & Baptista, A. M. (2004). A comparison of var and cvar constraints on portfolio selection with the mean-variance model. *Management science*, *50*(9), 1261–1273.

Ang, A., & Bekaert, G. (2004). How regimes affect asset allocation. *Financial Analysts Journal*, *60*(2), 86–99.

Ang, A., & Timmermann, A. (2012). Regime changes and financial markets. *Annu. Rev. Financ. Econ.*, *4*(1), 313–337.

Araújo, B. V. F., Camargos, M. A. d., & Pinho, F. M. d. (2018). Modeling conditional volatility by incorporating non-regular trading hours into the aparch model. *Revista Contabilidade & Finanças*, *30*, 202–215.

Ardia, D., Bluteau, K., & Rüede, M. (2019). Regime changes in bitcoin garch volatility dynamics. *Finance Research Letters*, *29*, 266–271.

Azad, N. F., & Serletis, A. (2022). Spillovers of us monetary policy uncertainty on inflation targeting emerging economies. *Emerging Markets Review*, *51*, 100875.

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

Black, F., & Litterman, R. (1992). Global portfolio optimization. *Financial analysts journal*, *48*(5), 28–43.

Borio, C. (2014). The financial cycle and macroeconomics: What have we learnt? *Journal of banking & finance*, *45*, 182–198.

Boyd, S., Busseti, E., Diamond, S., Kahn, R. N., Koh, K., Nystrup, P., Speth, J., et al. (2017). Multi-period trading via convex optimization. *Foundations and Trends® in Optimization*, *3*(1), 1–76.

Campani, C. H., & Garcia, R. (2019). Approximate analytical solutions for consumption/investment problems under recursive utility and finite horizon. *The North American Journal of Economics and Finance*, *48*, 364–384.

Campani, C. H., Garcia, R., & Lewin, M. (2021). Optimal portfolio strategies in the presence of regimes in asset returns. *Journal of Banking & Finance*, *123*, 106030.

Chekhlov, A., Uryasev, S., & Zabarankin, M. (2005). Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance*, *8*(01), 13–58.

Chopra, V. K., & Ziemba, W. T. (2013). The effect of errors in means, variances, and covariances on optimal portfolio choice. In *Handbook of the fundamentals of financial decision making: Part i* (pp. 365–373). World Scientific.

Claessens, S., Kose, M. A., & Terrones, M. E. (2012). How do business and financial cycles interact? *Journal of International economics*, *87*(1), 178–190.

Cui, X.-Y., Gao, J.-J., Li, X., & Shi, Y. (2022). Survey on multi-period mean–variance portfolio selection model. *Journal of the operations research society of china*, *10*(3), 599–622.

Daniel, P. P. (2025). *Portfolio optimization: Theory and application*. CAMBRIDGE University Press.

De Jong, M. (2018). Portfolio optimisation in an uncertain world. *Journal of Asset Management*, *19*(4), 216–221.

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The review of Financial studies*, *22*(5), 1915.

DeMiguel, V., Martín-utrera, A., & Uppal, R. (2024). A multifactor perspective on volatility-managed portfolios. *The Journal of Finance*, *79*(6), 3859–3891.

Ding, R. (2023). F-betas and portfolio optimization with f-divergence induced risk measures. *Quantitative Finance*, *23*(10), 1483–1496.

Gilbert, E., & Meiklejohn, L. (2019). A comparative analysis of risk measures: A portfolio optimisation approach. *Investment Analysts Journal*, *48*(3), 223–239.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the econometric society*, 357–384.

Han, C. (2020). How much should portfolios shrink? *Financial Management*, *49*(3), 707–740.

Harris, R. D., & Mazibas, M. (2013). Dynamic hedge fund portfolio construction: A semi-parametric approach. *Journal of Banking & Finance*, *37*(1), 139–149.

Hwang, I., Xu, S., & In, F. (2018). Naive versus optimal diversification: Tail risk and performance. *European Journal of Operational Research*, *265*(1), 372–388.

Israelsen, C. (2005). A refinement to the sharpe ratio and information ratio. *Journal of Asset Management*, *5*(6), 423–427.

Jarrow, R., & Li, S. (2021). Concavity, stochastic utility, and risk aversion. *Finance and Stochastics*, *25*, 311–330.

Kircher, F., & Rösch, D. (2021). A shrinkage approach for sharpe ratio optimal portfolios with estimation risks. *Journal of Banking & Finance*, *133*, 106281.

Kolm, P. N., Tütüncü, R., & Fabozzi, F. J. (2014). 60 years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research*, *234*(2), 356.

Kritzman, M., Page, S., & Turkington, D. (2012). Regime shifts: Implications for dynamic strategies (corrected). *Financial analysts journal*, *68*(3), 22–39.

Krokhmal, P., Uryasev, S., & Zrazhevsky, G. (2005). Numerical comparison of conditional value-at-risk and conditional drawdown-at-risk approaches: Application to hedge funds. In *Applications of stochastic programming* (pp. 609–631). SIAM.

Li, X., Uysal, A. S., & Mulvey, J. M. (2022). Multi-period portfolio optimization using model predictive control with mean-variance and risk parity frameworks. *European Journal of Operational Research*, *299*(3), 1158–1176.

Lopez-Ibanez, M., Dubois-Lacoste, J., Stützle, T., & Birattari, M. (2016). The irace package: Iterated racing for automatic algorithm configuration. *Operations Research Perspectives*, *3*, 43–58.

Lorimer, D. A., van Schalkwyk, C. H., & Szczygielski, J. J. (2024). Portfolio optimisation using alternative risk measures. *Finance Research Letters*, *67*, 105758.

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

Michaud, R. O. (1989). The markowitz optimization enigma: Is 'optimized'optimal? *Financial analysts journal*, *45*(1), 31–42.

Nystrup, P., Boyd, S., Lindström, E., & Madsen, H. (2019). Multi-period portfolio selection with drawdown control. *Annals of Operations Research*, *282*(1), 245–271.

Nystrup, P., Hansen, B. W., Madsen, H., & Lindström, E. (2015). Regime-based versus static asset allocation: Letting the data speak. *The Journal of Portfolio Management*, 103.

Nystrup, P., Madsen, H., & Lindström, E. (2018). Dynamic portfolio optimization across hidden market regimes. *Quantitative Finance*, *18*(1), 83–95.

Oprisor, R., & Kwon, R. (2021). Multi-period portfolio optimization with investor views under regime switching. *Journal of Risk and Financial Management*, *14*(1).

Ortobelli, S., Rachev, S. T., Stoyanov, S., Fabozzi, F. J., & Biglova, A. (2005). The proper use of risk measures in portfolio theory. *International Journal of Theoretical and Applied Finance*, *8*(08), 1107–1133.

Peng, C., Kim, Y. S., & Mittnik, S. (2022). Portfolio optimization on multivariate regime-switching garch model with normal tempered stable innovation. *Journal of Risk and Financial Management*, *15*(5), 230.

Qian, E. (2005). Risk parity portfolios: Efficient portfolios through true diversification. *Panagora Asset Management*.

Rockafellar, R. T., Uryasev, S., et al. (2000). Optimization of conditional value-at-risk. *Journal of risk*, *2*, 21–42.

Roman, D., Mitra, G., & Spagnolo, N. (2010). Hidden markov models for financial optimization problems. *IMA Journal of Management Mathematics*, *21*(2), 111–129.

Staden, P. M., Dang, D.-M., & Forsyth, P. A. (2021). The surprising robustness of dynamic mean-variance portfolio optimization to model misspecification errors. *European Journal of Operational Research*, *289*(2), 774–792.

Tu, J. (2010). Is regime switching in stock returns important in portfolio decisions? *Management Science*, *56*(7), 1198–1215.

Wobbrock, J. O., Findlater, L., Gergle, D., & Higgins, J. J. (2011). The aligned rank transform for nonparametric factorial analyses using only anova procedures. *Proceedings of the SIGCHI conference on human factors in computing systems*, 143–146.

Zhu, S., Zhu, W., Pei, X., & Cui, X. (2020). Hedging crash risk in optimal portfolio selection. *Journal of Banking & Finance*, *119*, 105905.