Portfolio Allocation under Sovereign Risk: Evidence from Brazil

Maria Ferreira[†, a]
Marcelo Fernandes[‡, a]

[a]Sao Paulo School of Economics, FGV

Abstract This paper examines a portfolio allocation strategy in which investment decisions depend on whether sovereign risk is high or low. We gauge sovereign risk by means of the spread on the 5-year credit default swap (CDS). The central premise is that the CDS spread provides a timely, market-based proxy for institutional risk, enabling portfolio weights to adjust dynamically across low- and high-risk environments. We show that the mean-variance portfolio conditional on the CDS spread regime entails a better risk-adjusted performance in the Brazilian market than the unconditional mean-variance portfolio and equal-weighted portfolio. This is particularly true for periods of high sovereign risk.

Keywords: asset allocation; CDS spread; mean-variance optimization; risk premium.

JEL codes: G11, C58.

1. Introduction

Portfolio allocation lies at the core of asset pricing since Markowitz's (1952) mean–variance framework. The efficient mean-variance frontier consists of portfolios that achieve the highest expected return for a given level of risk or, equivalently, the lowest risk for a given target return. Although conceptually parsimonious, the mean–variance approach remains influential, serving as the foundation for several approaches, including factor-based strategies and risk-oriented allocations such as risk budgeting and risk parity (Brandt, 2010; Roncalli, 2013).

In emerging markets, such as Brazil, abrupt shifts in the macroeconomic and institutional environment are frequent enough to challenge any assumption about time-invariant expected returns and covariance matrix. Recurrent

How to cite: Ferreira, M. & Fernandes, M. (2026). Portfolio Allocation under Sovereign Risk: Evidence from Brazil. *Revista Brasileira de Finanças*, *X*, e202601. https://doi.org/10.xxxx/xxxx
Submitted on January 1, 2026. Published online on March 31, 2026.
Editor in charge: Editor.

† mariasouzaf@icloud.com
‡ marcelo.fernandes@fgv.br

episodes of stress, stemming from fiscal imbalances, political uncertainty, and global financial shocks, often lead to pronounced changes in the distributions of asset returns, thereby undermining sample estimates of expected returns and their covariances. Under these circumstances, we argue that sovereign risk plays a central role. Sovereign CDS spreads provide a high-frequency, market-based measure of country risk that encapsulates investors' collective assessment of fiscal solvency and institutional credibility.

The literature widely documents that asset returns and covariances change across market states. During crisis periods, volatility and correlations typically rise, whereas they subside in more favorable market conditions (Ang and Bekaert, 2002). Regime-switching models provide a natural framework to capture such a nonstationary behavior in view that they allow the distribution of asset returns to depend on the prevailing market state (Hamilton, 1989). The empirical evidence for a wide array of financial markets, including Brazil, suggests that conditional allocation strategies tend to outperform static approaches, particularly if allowing for reallocation across asset classes in response to changes in the risk environment (Ang and Bekaert, 2002; Lewin and Campani, 2020; Pereira and Oliveira, 2021).

We contribute to this debate by proposing a portfolio-allocation strategy that conditions on the current regime (high or low) of the sovereign risk, as measured by Brazil's 5-year CDS spread. Unlike approaches based on latent regimes, we adopt an observable indicator to partition the sample period into times of low and high risk. Our key contribution is to employ the sovereign CDS spread as an explicit regime classifier for the estimation of conditional moments within the portfolio optimization process.

We evaluate our strategy under realistic operational constraints relative to standard benchmarks in the literature, including the unconditional mean-variance portfolio and the equal-weighted portfolio. The empirical results indicate that our conditional approach yields better risk-adjusted performance than both benchmarks, exhibiting not only higher Sharpe ratios but also lower drawdowns over the sample period. Portfolio decisions conditional on the CDS spread regime allow for reallocations that are more consistent with the prevailing risk environment, thereby enhancing diversification efficiency. Altogether, by integrating a market-based sovereign risk indicator directly into the portfolio optimization process, we contribute to the Brazilian asset allocation literature. In particular, we offer empirical evidence on the operational role of CDS spreads as a tool for dynamic portfolio management in emerging markets.

The remainder of the paper proceeds as follows. Section 2 reviews the literature on portfolio optimization, paying particular attention to regime-switching approaches. Section 3 describes the data, the regime-classification procedure, and the optimization methodology. Section 4 discusses the empirical results, whereas Section 5 reports some robustness checks. Section 6 offers some concluding remarks, including directions for future research.

2. Related literature

The canonical mean-variance problem defines the efficient frontier as the set of portfolios that minimize variance for each level of expected return or, equivalently, that maximize expected return for a given level of risk (Markowitz, 1952). Although short-selling constraints are not part of the original formulation, it is standard to impose them to put some discipline in the optimal weight estimation (Brandt, 2010). The resulting mean-variance problem remains a convex quadratic program, allowing for straightforward numerical optimization, and preserves the interpretation of efficiency under a restricted feasible set.

Tobin (1958) introduces the separation between the risky portfolio decision and the allocation between risky and risk-free assets, giving rise to the separation theorem. Sharpe (1964) and Lintner (1965) derive the capital asset pricing model (CAPM) from the mean–variance foundation, establishing a linear relationship between systematic risk and expected return. Later extensions, including (Black and Litterman, 1992), further refine the integration between equilibrium reasoning and portfolio construction. As such, the mean–variance paradigm remains the normative core of portfolio allocation.

Markowitz's original framework considers that expected returns and covariances are constant over time. This assumption is particularly unrealistic for emerging markets, where structural breaks, political shocks, fiscal crises, and external disturbances frequently affect the distribution of asset returns. Longin and Solnik (2001) and Ang and Bekaert (2002) show that the correlation between asset returns typically increases during periods of market stress, thereby reducing diversification precisely when it is most needed. In addition, they reveal that the distribution of asset returns is state-dependent and reacts asymmetrically across high- and low-risk environments.

The portfolio implications of regime dependence are substantial. Clarke and de Silva (1998) report that the existence of multiple states affects the shape of the efficient frontier, while Ramchand and Susmel (1998) identify regime-dependent volatility and correlation structures in Latin American equity markets, implying potential gains from conditional allocation. Ang and Bekaert (2002) further show that dynamic allocation strategies conditioned on regimes outperform static strategies, particularly when investors adjust exposures across equities, fixed income, and cash in response to state transitions. In the Brazilian context, Oliveira and Valls-Pereira (2018) demonstrate that entertaining multiple regimes significantly alters the composition of the tangency portfolio, whereas Lewin and Campani (2020) and Campani et al. (2021) find that tangency and risk-parity portfolios conditional on multiple regimes outperform their unconditional counterparts in terms of Sharpe ratio.

In this setting, sovereign risk emerges as a particularly relevant state variable. The sovereign CDS spread provides a high-frequency, forward-looking measure of country risk, incorporating fiscal, political, liquidity, and external vulnerability components. Unlike global volatility indicators such as the VIX, which primarily capture international risk sentiment, sovereign CDS spreads may reflect idiosyncratic domestic shocks. Episodes such as political crises, rating downgrades, and electoral uncertainty in Brazil trigger sharp increases in the CDS spread even if global volatility remains stable. Ma et al. (2018) interestingly show that CDS spreads exhibit regime-dependent dynamics, with flight-to-quality effects characterizing times of high sovereign risk. Consiglio et al. (2018) further demonstrate that CDS spreads are able to identify structural breaks and transitions throughout turbulence, crisis, and post-crisis phases.

These insights suggest that conditioning moments on sovereign-risk regimes could materially affect the optimal mean-variance portfolio. If return distributions differ across high- and low-risk states, using unconditional moments effectively distorts the estimation of the tangency portfolio. In contrast, conditional estimation allows portfolio weights to respond to shifts in expected returns and covariance structures as sovereign risk evolves.

The regime-dependent framework leads to three empirical predictions. First, if sovereign credit conditions contain incremental information about the distribution of asset returns, conditioning portfolio weights on the CDS spread should improve risk-adjusted performance relative to the unconditional mean–variance allocation. Second, by differentiating between high- and low-risk states, our approach should also alleviate downside risk. If the CDS effectively captures periods of elevated systemic stress, conditioning allocations on this signal should reduce the severity of drawdowns and improve tail-risk characteristics relative to unconditional benchmarks. Third, the benefits of regime conditioning should remain robust to reasonable variations in model design. If performance improvements arise from economically meaningful state differentiation, rather than overfitting, similar qualitative results should emerge under alternative CDS thresholds, volatility targets, and estimation-window lengths.

In general, while mean-variance theory provides the normative framework for constructing efficient portfolios, regime-dependent modeling recognizes that fiscal, political, and external shocks materially alter risk–return parameters in emerging markets. Despite the growing literature on regime-switching and emerging-market allocation, this is the first work to estimate the mean-variance portfolio conditional on CDS spread regimes within a realistic implementation setting.

3. Data and methodology

We consider both domestic equity and fixed-income indices in Brazil, as well as international equity and fixed-income assets. In particular, we use the Ibovespa as the Brazilian stock-market index, whereas the IMA-B and IMA-B5 indices proxy for inflation-linked government bonds and the CDI rate for the local risk-free rate. As for the international equity and fixed-income markets, we employ the S&P 500 index and the short-term T-bill index in the US. We convert the foreign indices into BRL using the BRL/USD exchange rate so as to take the perspective of local investors in Brazil.[1]

We focus on a small set of broad indices for two reasons. First, it reflects common practice among many investors, who often hold portfolios with a limited number of core positions. Second, it improves numerical stability in mean–variance optimization, which is very sensitive to estimation error if the number of assets is large relative to the number of time-series observations (Brandt, 2010).

We gather data from Bloomberg, assembling daily total returns for every index from January 2005 to December 2024, expressed in local currency (BRL). We use the first five years of the sample exclusively to initialize model estimates, so that the out-of-sample implementation begins in 2010.

3.1 Sovereign-risk regimes based on CDS spreads

The key contribution of this study is to condition portfolio allocation on risk regimes identified through the sovereign credit market. A sovereign CDS is a derivative instrument that reflects the cost of protection against a sovereign default. The CDS spread, as quoted in basis points (bps), rises when perceived country risk increases and declines during periods of higher confidence. As such, it serves as a daily measure of sovereign risk, incorporating both domestic factors (e.g., fiscal and political conditions) and external drivers (e.g., global risk appetite and international liquidity).

[1] Alternatively, one could use exchange-traded funds that track these indices at B3.

We use the Brazilian 5-year CDS spread as a country-risk signal that guides optimal portfolio allocation. Figure A1 in Appendix A shows that the CDS spread spikes coincide with major episodes of economic and political stress, such as the 2008-2009 global financial crisis, the 2015-2016 recession and impeachment period in Brazil, and the COVID-19 shock in 2020. This evinces the indicator's sensitivity to sovereign risk perceptions and global risk aversion. We also observe pronounced increases around the 2018 and 2022 elections, reflecting well the risk premia demanded by investors due to the high political uncertainty.

To reduce noise without delaying economically relevant information, we smooth the daily CDS spread using a 5-day moving average. Portfolio decisions at date t consider information available up to $t-1$ (that is to say, the CDS spread of the previous trading day). This is to prevent look-ahead bias. Using CDS as a state variable has the additional advantage of summarizing, at daily frequency, both domestic and global shocks that jointly affect the risk premia in Brazilian markets.

We classify high- and low-risk regimes in a dynamic and recursive manner. At each rebalancing date, we compute using a 5-year rolling window the top tercile of the empirical distribution of the daily smoothed CDS spread. If the CDS spread at $t-1$ exceeds the second tercile (i.e., 66th percentile), we then classify the period as high risk, otherwise as low risk. The threshold adapts to the prevailing risk level over time, preserving historical comparability without imposing a single fixed cutoff for the entire sample. Conditional on the prevailing regime, we estimate the expected returns and covariance matrix using only past observations from the same regime. This approach allows the portfolio to respond to regime changes without look-ahead bias, apart from capturing in a disciplined way the asymmetry between low and high sovereign-risk environments.

The choice of the top tercile of the CDS spread distribution as the regime determinant reflects a trade-off between statistical stability and economic relevance. On the one hand, stress regimes should not be so rare as to prevent reliable estimation of the expected returns and covariance matrix. On the other hand, the cutoff must be sufficiently high to capture episodes in which sovereign risk materially deviates from prevailing levels. In practical terms, defining high-risk regimes as observations in the upper third of the empirical distribution ensures an adequate number of data points for stable estimation while preserving the economic interpretation of stress periods. Additional analyses indicate that a lower cutoff at the median dilutes the notion of stress, whereas a higher cutoff at 80th percentile yields excessively sparse high-risk regimes, thereby reducing statistical power.

The resulting regime indicator then reads

$$z_{t-1} = \mathbf{1}\left\{\frac{1}{5}\sum_{j=1}^{5}\mathrm{CDS}_{t-j} > \tau_t^{(2/3)}\right\}, \tag{1}$$

where $\tau_t^{(2/3)}$ is the second tercile of the empirical distribution of the smoothed CDS spreads at time t. For each 5-year rolling window at time t, we denote the dates at which the regime indicator is equal to one as the subsample of high-risk periods $\mathscr{H}_t$, whereas the dates at which $z_{t-1} = 0$ form the subsample of low-risk periods $\mathscr{L}_t$. We update this classification at each rebalancing date t, so as to track the evolution of sovereign risk over time without imposing a fixed threshold for the entire sample.

We then carry out the estimation of the inputs in the mean-variance portfolio optimization by conditioning on the prevailing regime. Specifically, we consider only the subset of observations within the 5-year rolling window that belongs to the same regime as $t-1$. This means that, if the regime is of high risk at $t-1$, we estimate the expected returns and covariance matrix using only historical observations in $\mathscr{H}_t$. Analogously, we employ only observations in $\mathscr{L}_t$ to estimate the moments of the distribution if the regime at $t-1$ is of low risk.

Based on the conditional subsamples, we estimate expected returns by the vector $\hat{\mu}_{\mathscr{I}_t}$ of average returns over the days in $\mathscr{I}_t \in \{\mathscr{H}_t, \mathscr{L}_t\}$. As for the conditional covariance matrix, we employ the exponentially weighted moving average (EWMA) estimator, with $\lambda = 0.94$ as in Suganuma (2000), but conditional on the CDS spread regime. This means we assign more weight to recent observations in that

$$\widehat{\Sigma}_{\mathscr{I}_t} = \frac{1}{\#\mathscr{I}_t}\sum_{s\in\mathscr{I}_t}\frac{\lambda^{(T-s)}}{\sum_{\tau\in\mathscr{I}_t}\lambda^{(T-\tau)}}(r_s - \hat{\mu}_{\mathscr{I}_t})(r_s - \hat{\mu}_{\mathscr{I}_t})^\top, \tag{2}$$

where T is the most recent observation in the subsample $\mathscr{I}_t$, and the decay factor $\lambda \in (0,1)$ controls the rate at which past information loses relevance.

The expected return in excess over the CDI risk-free rate then is

$$\hat{\mu}^e_{\mathscr{I}_t} = \frac{1}{\#\mathscr{I}_t}\sum_{s\in\mathscr{I}_t}(r_s - \mathrm{CDI}_s), \tag{3}$$

which drives our minimum-return constraint in the optimization problem by imposing that $w_t^\top \hat{\mu}_t^e \geqslant 0$. The latter establishes a floor on the *ex-ante* expected risk premium over the risk-free proxy, ruling out portfolios that would underperform the CDI rate on average. It is important to stress that this amounts to an *in-sample* condition based on training-window estimates, and hence *ex-post* realized returns may still fall below the CDI rate due to sampling error. Analogously, even if we entertain a target for the *ex-ante* portfolio volatility, the realized volatility may differ from the target volatility because of sampling error, as well as regime shifts outside the estimation window.

Overall, our allocation strategy is quite simple, relying exclusively on past information (i.e., no look-ahead bias). The input estimates reflect the statistical behavior we observe in the historical subsamples within each regime. As a result, conditional estimates capture heterogeneity in returns and dependence structures across sovereign-risk states, enabling portfolio choices that are more responsive to the prevailing macro-financial environment.

3.2 Portfolio-allocation procedures

This section outlines the three portfolio allocation strategies we evaluate in our empirical analysis. Apart from the equal-weighted benchmark (1/N), we also compute the conditional and unconditional mean-variance portfolios based on expected returns and covariance matrix estimates coming either from the subsample based on the current CDS spread regime or from the entire sample, respectively.

The equal-weighted strategy assigns identical weights to each of the N assets in the investment universe. Although simplistic, it is not straightforward to outperform the 1/N portfolio because it achieves substantial diversification without requiring any parameter estimation or model assumptions (DeMiguel et al., 2009). It provides a useful reference to assess whether more sophisticated optimization techniques effectively add value. In the present setting, it produces a balanced allocation across asset classes but does not necessarily satisfy the expected return and volatility targets.

The unconditional mean-variance (UMV) portfolio corresponds to the standard mean–variance optimization without accounting for the CDS spread regime. It corresponds to the perspective of the investor that ignores the information content of the CDS spread. Accordingly, we estimate the expected returns and covariance matrix using every observation in the rolling window of the most recent five years of daily data, regardless of the risk state prevailing during that period. This specification serves as a control model, isolating the incremental contribution of conditioning on high- and low-risk regimes.

The conditional mean-variance (CMV) strategy explicitly accounts for the sovereign-risk regime as identified by the CDS spread signal in the previous day. At each rebalancing date, if the market is in the low-risk regime, we estimate the expected returns and covariance matrix using only the subsample $\mathscr{L}_t$. Analogously, if the current regime is of high sovereign risk, we estimate the parameters using exclusively the observations in the subsample $\mathscr{H}_t$. Conceptually, this approach generates two distinct efficient frontiers: namely, one corresponding to normal market conditions and another in times of market stress. The portfolio dynamically switches between these frontiers according to the prevailing regime. By conditioning the moment estimation on the current risk state, the strategy aims to produce *ex-ante* allocations that are more closely in line with the prevailing market conditions.

We implement every strategy over the same universe of $N = 6$ assets: namely, Ibovespa, IMA-B, IMA-B5, CDI, S&P 500 index and T-bill. They are also subject to identical weight constraints, ensuring full comparability across specifications: portfolios are fully invested, with neither leverage nor short selling such that individual weights are non-negative and sum up to one. Rebalancing takes place every 15 trading days to alleviate transaction costs. We impose an *ex-ante* volatility target of 10% per year to act as a risk cap. Additionally, as aforementioned, we require portfolios to satisfy a non-negative expected return in excess over the CDI rate. Finally, we penalize portfolio returns *ex-post* by means of a transaction cost of 10 basis points per unit of turnover before computing realized performance.

3.3 Parameter estimation and portfolio optimization

Let $r_s \in \mathbb{R}^N$ denote the vector of asset returns at day s. At each rebalancing date t, we identify the prevailing regime $z_{t-1} \in \{0,1\}$ by checking whether the CDS spread belongs to the top tercile: if yes, $z_{t-1} = 1$ and $\mathscr{I}_t = \mathscr{H}_t$; if not $z_{t-1} = 0$ and $\mathscr{I}_t = \mathscr{L}_t$. It may happen that there are not enough observations within the 5-year rolling window for estimating the expected returns and covariance matrix in a precise manner for both regimes. Accordingly, we impose a minimum number of 252 daily observations in each regime, expanding our rolling window to the past up to the date we meet this criterion. This procedure preserves the absence of look-ahead bias, while ensuring sufficient data for reliable moment estimation.

Given the conditional and unconditional expected returns and covariance matrix estimates, we obtain the optimal portfolio weights by solving the constrained mean-variance optimization problem with both volatility and excess return targets (i.e., risk cap and non-negative expected returns in excess over the CDI rate). The optimization problem at each rebalancing date t is as follows:

$$\begin{aligned}\max_{w\in[0,1]^N} \quad & w^\top \hat{\mu}_t \\ \text{s.t.} \quad & w^\top \widehat{\Sigma}_t w \leqslant \bar{\sigma}, \\ & \mathbf{1}^\top w = 1, \\ & w^\top \hat{\mu}_t^e \geqslant 0.\end{aligned} \tag{4}$$

where $\bar{\sigma}$ denotes the annual volatility target, which we set to 10% per year (conversion assumes 252 trading days per year).

The optimization problem in (4) seeks to maximize the expected portfolio return, but with constraints. The first imposes an *ex-ante* risk cap, shrinking the portfolio volatility to avoid exceeding the annual volatility target. The second constraint enforces full investment, while the third imposes a non-negative expected excess return relative to the risk-free rate. This framework ensures strict comparability across the mean-variance allocation strategies, which differ solely in the estimation of expected returns and covariance matrix. There is no difference with respect to constraints, volatility targets, and rebalancing frequency. The regime-dependent specification therefore incorporates sovereign risk information through conditional moment estimation, while the UMV and 1/N portfolios serve as unconditional benchmarks.

4. Does accounting for sovereign risk pay off in Brazil?

In this section, we document that conditioning portfolio optimization on sovereign risk regimes leads to economically meaningful differences in both asset allocation and risk–return characteristics. In particular, our empirical evidence supports our main hypothesis that, by segmenting market states through sovereign CDS information, investors can effectively reconstruct state-contingent efficient frontiers. This approach reduces the likelihood of excessive pro-cyclical exposure during adverse scenarios, thereby improving the overall risk–return trade-off of the portfolio.

4.1 Aggregate portfolio performance

Table 1 reports the aggregate performance of each allocation strategy (1/N, UMV and CMV) over the post-estimation sample that runs from 2010 to 2024. We compute the annualized return, annualized realized volatility, Sharpe ratio, Sortino ratio, maximum drawdown, and maximum/minimum daily returns of each strategy after transaction costs (i.e., net-of-transaction-cost wealth indices).

Table 1
Annualized performance net of transaction costs, 2010 to 2024

metric	CMV	UMV	1/N
return (% p.a.)	15.84	13.76	10.75
realized volatility (% p.a.)	10.43	10.86	6.01
Sharpe ratio	0.51	0.33	0.08
Sortino ratio	0.74	0.47	0.12
maximum drawdown (%)	-16.14	-19.30	-13.89
maximum daily return (%)	6.58	6.63	5.03
minimum daily return (%)	-3.76	-5.68	-5.63

Notes: The performance measures derive from net wealth indices over the out-of-sample period. Annualization assumes 252 trading days per year.

Our findings indicate that the CMV portfolio delivers the strongest risk-adjusted performance, yielding a return of 15.84% per year for a realized volatility very close to 10% per year. As such, Sharpe and Sortino ratios exceed those of the UMV and 1/N portfolios. From an economic standpoint, these results suggest that accounting for sovereign risk regimes enhances the alignment between portfolio exposures and the prevailing state-dependent risk structure. By separating high- and low-risk states, it dynamically adjusts allocation to reflect regime-specific expected return and covariance structures. Accordingly, it ends up capturing risk premia more efficiently during stable periods, while reducing vulnerability to adverse shocks when sovereign risk intensifies.

The UMV portfolio delivers the second best performance, revealing the gains of efficient diversification relative to the naive diversification of the 1/N portfolio. The UMV portfolio yields a competitive return of 13.76% per year, with a slightly higher realized volatility than the CMV portfolio. Apart from the lower Sharpe and Sortino ratios, the UMV portfolio also exhibits deeper drawdowns. This pattern is consistent with the use of unconditional moments that aggregate heterogeneous market states within a single estimation window. By averaging across distinct risk environments, the unconditional approach dilutes the regime-specific information, responding more gradually to structural shifts in sovereign-risk conditions.

The equal-weighted portfolio yields not only lower returns, but also lower realized volatility and maximum drawdown. The latter is not enough to compensate the drop in the average returns, though. As such, both Sharpe and Sortino ratios are very close to zero, despite the minimal rebalancing requirements. Interestingly, the realized volatility of 6.01% per year reflects the absence of explicit volatility or risk-budgeting controls. As a result, the naive diversification of the 1/N allocation offers a conservative benchmark that obviously does not adjust exposures in response to changing risk conditions.

Figure 1 displays how the net-of-transaction-cost indices evolve over time for each allocation rule from 2010 to 2024. While they follow similar trajectories in the early part of the sample, the regime-dependent portfolio gradually stands out from the other portfolios, particularly during episodes of heightened domestic stress. In these periods, the CMV portfolio indeed has relatively smaller drawdowns and less volatility. This pattern is consistent with the notion that conditioning on sovereign risk states improves the intertemporal consistency of portfolio risk exposures. Overall, the evidence suggests that controlling for the CDS spread regime enhances portfolio efficiency in the Brazilian market.

Figure 1
Performance of each allocation strategy from 2010 to 2024

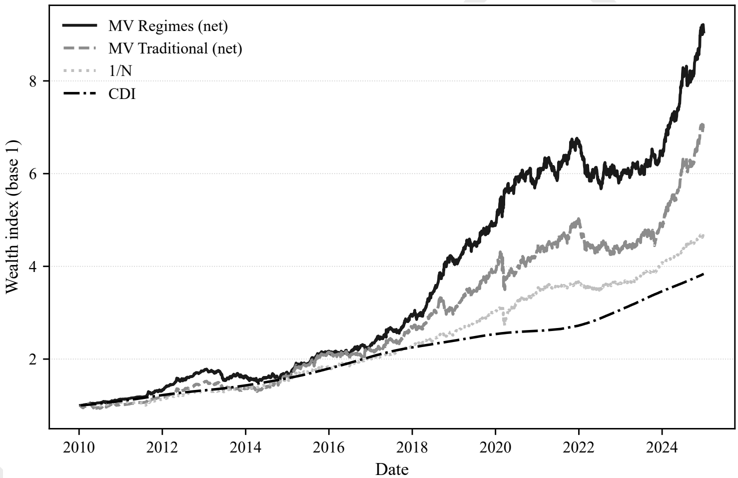

It is also interesting to observe that every strategy delivers returns well above the CDI rate, albeit with clear differences in diversification efficiency. The CMV portfolio exhibits shallower drawdowns during major stress episodes, such as the 2015–2016 political crisis and the 2020 pandemic shock. Figure 2 further shows that it displays a faster recovery after times of high uncertainty, reinforcing the notion that accounting for the sovereign risk enables not only the identification of shifts in the market conditions but also the necessary adjustment in exposures. Notably, the CMV strategy outperforms in the vast majority of stress-related calendar years, suggesting that the informational content of the CDS spread signal is indeed economically meaningful and hence not a statistical artifact. Figure 2 breaks down performance per year to offer some additional insight. It is patent that the conditional allocation rule seemingly outperforms more substantially in periods of macroeconomic and political uncertainty, such as 2018 and 2020, when sovereign risk conditions shift materially.

In contrast, in periods of relative macro-financial stability as in Figure 3, the difference in performance between the CMV and UMV portfolios narrows. Although the unconditional mean-variance portfolio might occasionally outperform the conditional portfolio in calm years, such as 2014 and 2019, the latter typically improves downside resilience without sacrificing much performance in normal times.

Figure 2
Performance at times of market stress in Brazil

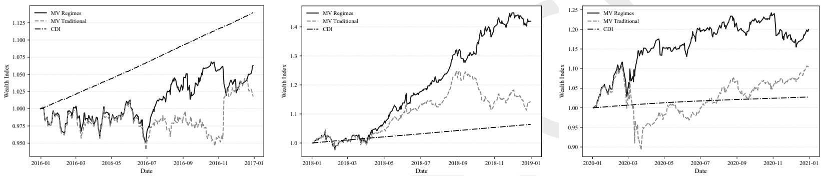

(a) recession + impeachment, 2016 **(b)** trucker's strike + elections, 2018 **(c)** COVID pandemic, 2020

Figure 3
Performance at times of relative tranquility in Brazil

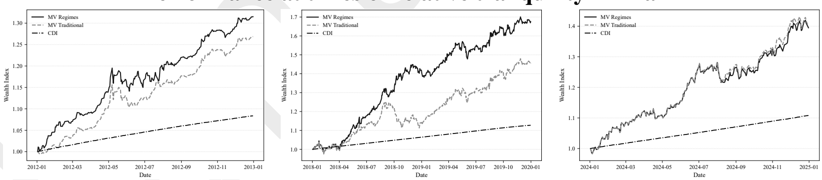

(a) low rates + liquidity, 2012 **(b)** post-impeachment, 2018—2019 **(c)** decrease in risk perception, 2024

From a practical standpoint, investors adopting the CMV strategy would have experienced more moderate capital losses and a smoother wealth trajectory over time. Figure 4 presents the drawdown profiles of each allocation strategy. The CMV portfolio exhibits moderately smaller peak losses and, in some stress episodes, a faster recovery than the UMV portfolio. The maximum drawdown of the conditional approach reaches -16.1%, comparing well to the maximum drawdown of -19.3% of the UMV portfolio.

While the difference in magnitude is not so substantial, it remains economically relevant in a portfolio-management context, as incremental reductions

Figure 4
Drawdown dynamics across portfolio strategies, 2010 to 2024

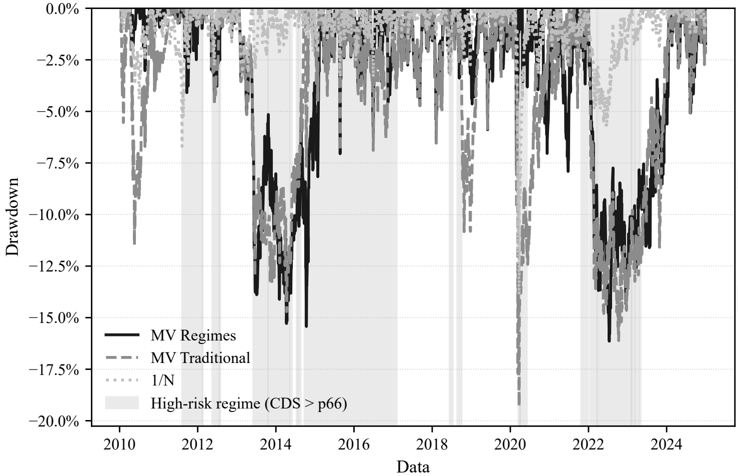

in drawdown depth compound meaningfully over longer horizons. As aforementioned, the two-regime approach enhances downside resilience without compromising return generation in normal periods. As such, it materially improves the risk-adjusted performance by mitigating loss episodes through sovereign-risk segmentation. Altogether, the results corroborate our core hypothesis: namely, that incorporating sovereign risk regimes into the portfolio optimization process enhances diversification efficiency both in the central and in the tail of the return distribution.

4.2 Allocation dynamics and turnover

4.2.1 Portfolio weights and sensitivity to sovereign-risk regimes

Figure 5 depicts how the optimal mean-variance portfolio weights evolve over time. Visual comparison indicates that the conditional specification implements substantial adjustments in portfolio composition as the Brazil 5-year CDS shifts between low and high sovereign-risk states. In contrast, the unconditional approach maintains a relatively stable structure over time, with more gradual variations and limited sensitivity to changes in the risk environment.

In low-risk conditions, the CMV portfolio allocates predominantly to long-duration inflation-linked Brazilian bonds (IMA-B), which represent approximately 40% of the portfolio on average. In addition, it recommends a substantial allocation to international assets: namely, around 35% in the S&P

Figure 5
Optimal portfolio weights under conditional and unconditional allocation rules

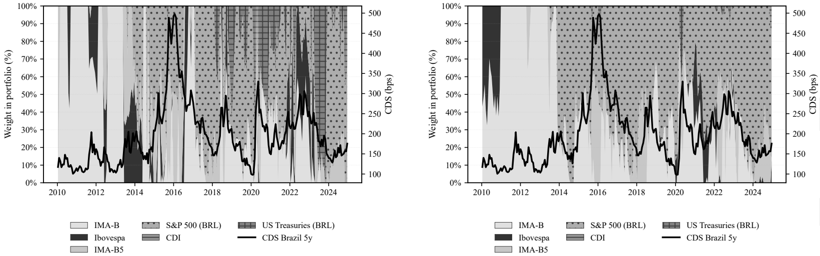

(a) conditional weights given CDS spread regime **(b)** unconditional mean-variance portfolio weights

500 index and 16% in the US T-bills, both in BRL terms. This configuration is consistent with a benign sovereign-risk scenario in which domestic real yields remain attractive and it is possible to pursue international diversification without breaching the volatility target. The combination of long-duration local bonds and global assets suggests that, in low-risk regimes, the CMV portfolio exploits both domestic term premia and cross-border diversification benefits.

In the high-risk regime, allocation changes dramatically in the CMV portfolio. Exposure to shorter-duration inflation-linked bonds (IMA-B5) increases from roughly 7% to 19%, implying a major reduction in duration risk. More strikingly, the average allocation to Ibovespa increases sharply, from approximately 0.8% in low-risk states to more than 25% under high-risk conditions. This shift represents the most pronounced regime-driven adjustment in the CMV portfolio composition.

Rather than adopting a purely defensive posture, the two-regime strategy reconfigures the exposure to domestic risk. The increase in Ibovespa allocation during periods of high sovereign risk may appear counterintuitive at first glance. Conventional wisdom would likely suggest a reduction in local equity exposure in times of high sovereign risk. However, the CDI rate is typically high in such periods, raising the bar implied by the constraint on the portfolio's excess return. This means the CMV portfolio will have to reallocate weight toward assets that help satisfy the excess-return constraint at the lowest variance cost. Table A1 in Appendix A shows that this equates to a materially higher exposure to domestic equities during stress episodes given that international assets exhibit relatively lower average returns in BRL. This pattern does not reflect mechanical risk-taking behavior; instead, it is the outcome of a state-contingent re-optimization under a different risk-return profile.

In contrast, the UMV portfolio maintains a relatively stable allocation across the sample period, with a quite persistent concentration in IMA-B and in the S&P 500 index. Because unconditional moments average out the high- and low-risk scenarios, the resulting portfolio smooths out the regime-specific dynamics, ironing out shifts in domestic risk exposure when sovereign risk increases in a significant manner. Accordingly, incorporating sovereign-risk regimes into the mean-variance optimization program changes the cyclical behavior of portfolio allocation in a considerable manner. By allowing portfolio weights to respond endogenously to regime-specific moments, the CMV portfolio achieves greater resilience and tighter alignment with the structural features of the Brazilian markets, where both regime shifts and confidence shocks seem recurrent.

4.2.2 Turnover considerations

We next carry out a turnover analysis to examine how the conditional approach translates risk reassessments into portfolio adjustments. More precisely, we define turnover as the aggregate variation in portfolio weights at each rebalancing date. By comparing the weight updates of the CMV and UMV portfolios, it is possible to assess whether the informational gains derived from sovereign-risk segmentation persist after accounting for the implicit costs of reallocation.

As expected, we find that the CMV portfolio has a higher turnover than the UMV portfolio. The average turnover at a given rebalancing day is 44.5% in the CMV portfolio, against 25.9% in the unconditional strategy. This difference reflects the higher responsiveness of the conditional allocation rule to changes in the 5-year Brazil CDS spread and to shifts in regime-specific conditional moments. While the unconditional optimization gives way to relatively stable weights across subsequent rebalancing dates, the regime-driven allocation enforces sharper reallocations in periods of risk transition, thereby adapting more rapidly to changes in market conditions.

Despite the higher turnover, rebalancing costs remain under control in that the average rebalancing cost is of approximately 4.5 bps and 2.6 bps respectively for the CMV and UMV portfolios. In view that we rebalance the portfolios only once every 15 trading days, these figures imply an expected drag of about 0.75% and 0.43% per annum, respectively. From an economic perspective, conditioning on CDS spread regimes improves the risk-adjusted performance of the portfolio allocation at a quite modest cost.

Turnover spikes, as captured by the 95th percentile, further highlight the episodic nature of the changes in the CMV portfolio weights. Turnover levels reaches nearly 173% during stress events, whereas the UMV turnover never exceeds 85%. This suggests that we change the CMV portfolio weights very sharply in some very stressful moments (as measured by the CDS spread), such as political crises or liquidity shocks. This is exactly the pattern that arises in Figure 6, which portrays how the turnover of both portfolios varies over time. The conditional approach exhibits a clearly episodic pattern, with adjustments concentrated around transitions in sovereign risk states, whereas the UMV portfolio shows more gradual and dispersed adjustments.

Figure 6
Turnover sensitivity to CDS spread regimes

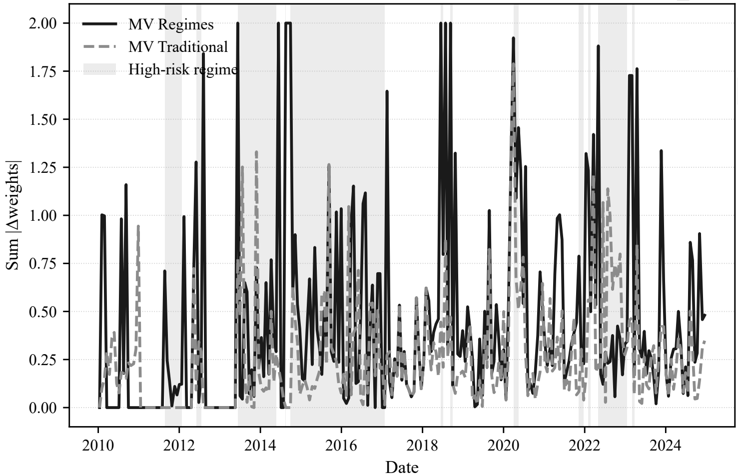

This behavior underscores the heightened sensitivity of the conditional allocation rule to regime shifts, due to the swifter updates in the expected returns and covariance matrix estimates. By reacting promptly to spikes in the CDS spread, the CMV strategy seeks to protect the portfolio by reducing duration and reallocating toward assets with lower conditional correlation, at the cost of higher turnover. In contrast, the UMV allocation rule adjusts weights more smoothly and with a larger delay.

5. Robustness checks

To assess whether the empirical gains of the allocation rule based on CDS spread regimes are genuine, we conduct robustness experiments along three dimensions: (*i*) the percentile of the CDS spread distribution that defines the high and low risk regimes, (*ii*) the target volatility in the optimization problem, and (*iii*) the length of the rolling estimation window. The investment universe, transaction costs, rebalancing frequency, and other portfolio constraints remain otherwise constant.

5.1 Regime classification

We examine how the choice of the percentile of the CDS spread distribution affects the performance of the regime-driven optimization problem. In the previous sections, we consider dates at which the CDS spread is in the top tercile as those of high sovereign risk. We now evaluate the percentiles 50%, 70% and 80% using the same implementation protocols as before.

Table A3 in Appendix A reports the performance measures for the different percentiles. Interestingly, we find that the performance of the CMV portfolio does not vary monotonically with the CDS spread cutoff in that the alternative percentiles yield lower average returns than the original. In fact, lower percentiles lead to a noisier high-risk classification, whereas larger percentiles increase the sampling error in the estimation of the expected returns and covariance matrix in the high-risk regime due to the smaller sample sizes.

The resulting CMV portfolios nonetheless remain broadly in line with the *ex-ante* target volatility. Maximum drawdowns become much deeper for more extreme percentiles, though. As the distribution of daily returns exhibits a good deal of skewness, the sample size of the conditional moment estimation contracts very quickly as we increase the percentile of the CDS spread distribution. This not only weakens the allocation responsiveness to adverse states, but also amplifies downside exposure. To sum up, the evidence suggests that using the top tercile offers the best tradeoff between economic differentiation of stress periods and stability of conditional moment estimates.

5.2 Volatility target

We next assess the sensitivity of the CMV portfolio to the volatility target. In a constrained mean–variance setting, the volatility target $\bar{\sigma}$ reflects the investor's risk budget. Lower targets prioritize capital preservation at the expense of return potential, whereas higher targets allow greater exposure to risk premia but increase vulnerability to market shocks and deeper drawdowns. We next examine how the portfolio allocation rules behave across different volatility targets to understand whether our previous findings are robust to varying levels of risk tolerance.

We consider a set of alternative volatility targets (namely, 5%, 8%, 15% and 20% per year), while holding constant every other implementation detail. To assess relative performance, we compute differences in performance with respect to the UMV portfolio. As before, we focus on average return, realized volatility, Sharpe and Sortino ratios, and maximum drawdown. Table A3 in Appendix A shows that the performance advantage of the regime-driven allocation increases with the volatility target. As the risk budget grows, the

CMV portfolio yields progressively larger excess returns relative to the UMV portfolio. In particular, the return differential rises steadily from below one percentage point at the smallest volatility level to more than three percentage points at the highest target.

Importantly, improvements are not confined to return alone. Across virtually all volatility targets, the realized volatility of the CMV portfolio remains slightly below that of the UMV portfolio, suggesting larger diversification gains. As for downside risk, maximum drawdowns are generally less severe for the CMV portfolios with conservative to intermediate volatility targets. If the risk budget is large enough, drawdowns become comparable to those of the UMV portfolios.

Because we obtain higher returns at similar levels of realized volatility, the CMV portfolio entails better risk-adjusted performance than the UMV portfolio across the entire spectrum of volatility targets. Sharpe and Sortino ratios remain consistently higher, despite monotonically declining as the volatility target increases. This pattern suggests that the benefits of conditioning on the CDS spread regime are not specific to a particular risk attitude. Instead, the informational content of the CDS signal enhances allocation efficiency for both conservative and aggressive investors.

5.3 Rolling-window length

We next examine the sensitivity to the length of the rolling estimation window we use to estimate the expected returns and covariance matrix in the different regimes. The estimation horizon plays a dual role. While it determines the effective sample size available to estimate the conditional moments within each regime, it also governs the trade-off between statistical stability and responsiveness to structural changes in the risk environment.

Short estimation windows increase adaptability to recent information, allowing the estimators to react more quickly to shifts in sovereign risk conditions. However, they also amplify estimation error, particularly in the high-risk regime, where the subsample size is substantially smaller by construction. Conversely, longer windows enhance statistical precision but could well dilute regime-specific dynamics by mixing heterogeneous states. To evaluate this trade-off, we revisit the empirical analysis using rolling windows of 3, 4, 5, 6, and 7 years to estimate the expected returns and covariance matrix. As before, we do not change any other implementation detail, and assess performance by computing the gains in average returns and Sharpe ratios relative to the unconditional allocation procedure.

Table A3 reveals a non-monotonic relationship between the length of the estimation window and relative performance. The CMV portfolio underperforms the UMV portfolio for 3-year rolling windows, which is consistent with noisier estimates. Relative performance remains close to neutral for 4-year estimation windows, but increases materially as we increase the window length to 5 or 6 years. Finally, we observe performance loss for excessively long windows of 7 years that dilute the informational content of the regime classification by blending structurally distinct risk states.

6. Conclusion

This study assesses the relative performance of a state-dependent portfolio allocation strategy for the Brazilian investor. We entertain two market states, reflecting either high or low sovereign risk. We proxy sovereign risk using Brazil's 5-year CDS spread and use data from 2010 to 2024 to construct the regime classification. We find that the optimal portfolio weights vary significantly depending on whether the CDS spread is in the high or low regime. Moreover, the conditional approach yields risk-adjusted performance gains relative to the unconditional optimization.

The key mechanism lies in the regime-specific reordering of relative asset efficiency. When sovereign risk is high, the conditional allocation favors mainly shorter-duration inflation-linked bonds and Ibovespa, whereas allocation shifts towards longer-duration and international assets in low-risk scenarios. These portfolio tilts emerge endogenously from changes in the conditional expected returns and covariances, rather than from *ad hoc* tactical rules.

Taken together, there are three main implications: (*i*) sovereign-risk regimes are informationally relevant in that the tangency portfolio depends on the risk regime, (*ii*) conditional allocation improves efficient diversification, exhibiting more resilient performance during stress episodes, and (*iii*) implementation is operationally feasible. Our findings reinforce the argument that portfolio management in Brazil requires explicit conditioning on macro-financial risk states. In environments characterized by recurrent political and fiscal shocks, unconditional moments obscure structural shifts and delay portfolio updates. A state-dependent allocation rule anchored in sovereign risk indicators provides a disciplined way to incorporate institutional risk into quantitative asset allocation.

We conclude by discussing two caveats. First, regime classification relies on historical information, assuming that the relationship between sovereign risk and asset returns is sufficiently persistent. Structural breaks in this relationship or shifts in the macro-financial transmission mechanism could reduce the predictive content of the CDS spread. Second, we carry out the empirical analysis with a specific emerging market in mind. The effectiveness of our state-dependent approach may vary across countries with different institutional, market and political conditions given that the CDS spread is just a catch-all measure of sovereign risk.

Acknowledgments Fernandes acknowledges financial support from FAPESP (2023/01728-0) and CNPq (302153/2022-5). The usual disclaimer applies.

Conflict of interest The authors declare no conflict of interest.

Artificial Intelligence This research utilized AI tools to assist in data analysis. The authors have critically reviewed and validated every AI-generated content to ensure accuracy and alignment with the scientific integrity of the study. The use of AI adhered to ethical guidelines, ensuring transparency and compliance with academic standards. The authors carefully considered any biases or limitations inherent to the AI tools in the interpretation of results. The authors affirm that the AI tools did not compromise the originality or integrity of the work.

Data availability Data are available from the corresponding author upon reasonable request.

References

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

Black, F. and Litterman, R. (1992). Global portfolio optimization, *Financial Analysts Journal*, *48*(5), 28–43.

Brandt, M. W. (2010). Portfolio choice problems, *Handbook of Financial Econometrics: Tools and Techniques*, Elsevier, pp. 269–336.

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

Clarke, R. G. and de Silva, H. (1998). State-dependent asset allocation, *Journal of Portfolio Management*, *24*(2), 57.

Consiglio, A., Lotfi, S. and Zenios, S. A. (2018). Portfolio diversification in the sovereign credit swap markets, *Annals of Operations Research*, *266*(1), 5–33.

DeMiguel, V., Garlappi, L. and Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?, *Review of Financial Studies*, *22*(5), 1915–1953.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica*, *57*(2), 357–384.

Lewin, M. and Campani, C. H. (2020). Portfolio management under multiple regimes: Strategies that outperform the market, *Revista de Administração Contemporânea*, *24*(4), 300–316.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics*, *47*(1), 13–37.

Longin, F. and Solnik, B. (2001). Extreme correlation of international equity markets, *Journal of Finance*, *56*(2), 649–676.

Ma, J. Z., Deng, X., Ho, K.-C. and Tsai, S.-B. (2018). Regime-switching determinants for spreads of emerging markets sovereign credit default swaps, *Sustainability*, *10*(8), 2730.

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

Oliveira, A. B. and Valls-Pereira, P. L. (2018). Asset allocation with Markovian regime switching: Efficient frontier and tangent portfolio with regime switching, *Brazilian Review of Econometrics*, *38*(1), 97–127.

Pereira, P. L. V. and Oliveira, A. B. (2021). Estratégias de investimento em portfólios com estimativas de alta e baixa do mercado financeiro, *Revista Brasileira de Finanças*, *19*(4), 160–185.

Ramchand, L. and Susmel, R. (1998). Volatility and cross correlation across major stock markets, *Journal of Empirical Finance*, *5*(4), 397–416.

Roncalli, T. (2013). *Introduction to Risk Parity and Budgeting*, CRC press.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance*, *19*(3), 425–442.

Suganuma, R. (2000). Reality check for volatility models, *University of California, San Diego*, .

Tobin, J. (1958). Liquidity preference as behavior towards risk, *Review of Economic Studies*, *25*(2), 65–86.

A. Supplementary Material

This section provides additional evidence on the construction of the sovereign-risk regimes and their empirical implications for asset returns.

Figure A1
Brazil's CDS spread and sovereign-risk regime classification

We display Brazil's five-year CDS spread and its second tercile considering a five-year window. The shaded areas correspond to the periods in which the CDS spread is in the top tercile and hence sovereign risk is high.

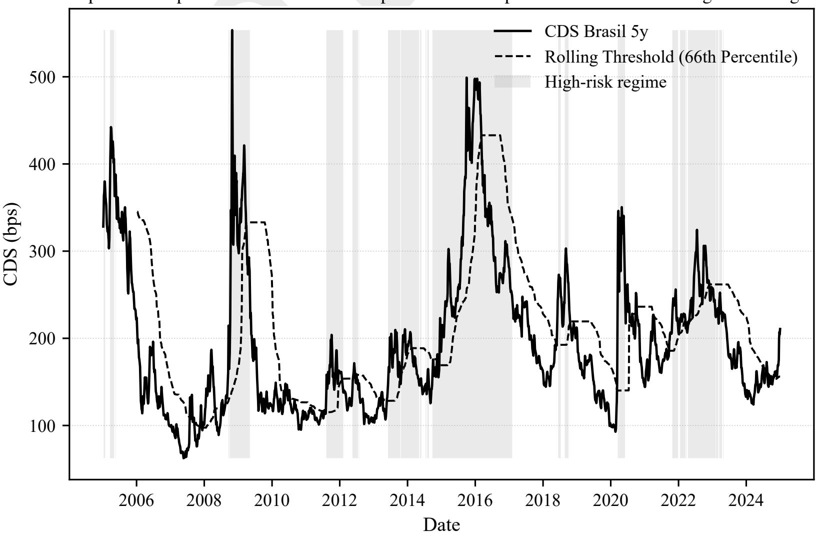

Table A1
Conditional asset performance by sovereign-risk regime

We report annualized average return (raw and in excess over the CDI rate, in BRL), realized volatility, Sharpe ratio, and the number of observations for each asset class, conditional on sovereign-risk regimes as defined by the rolling 66th percentile of the five-year CDS spread distribution.

asset	average return	excess return over CDI rate	realized volatility	Sharpe ratio	sample size
subsample of low sovereign risk					
CDI	9.35%	0.00%	0.24%	0.00	3338
IMA-B5	10.59%	1.24%	2.57%	0.48	3338
IMA-B	10.84%	1.49%	5.83%	0.26	3338
Ibovespa	5.41%	-3.94%	23.73%	-0.17	3338
S&P 500 (in BRL)	11.39%	2.04%	16.84%	0.12	3338
T-bill (in BRL)	6.42%	-2.93%	14.06%	-0.21	3338
subsample of high sovereign risk					
CDI	11.14%	0.00%	0.16%	0.00	1677
IMA-B5	13.02%	1.88%	3.31%	0.57	1677
IMA-B	12.89%	1.74%	8.06%	0.22	1677
Ibovespa	13.27%	2.13%	30.93%	0.07	1677
S&P 500 (in BRL)	13.75%	2.61%	23.87%	0.11	1677
T-bill (in BRL)	4.99%	-6.16%	20.67%	-0.30	1677

Table A2
Average weights of the CMV and UMV portfolios

We report average portfolio weights (in percent) for both CMV and UMV portfolios.

asset	low-risk regime	high-risk regime
state-dependent allocation rule, CMV portfolio		
CDI rate	0.40%	0.92%
IMA-B5	7.39%	19.19%
IMA-B	39.58%	26.45%
Ibovespa	0.82%	25.13%
S&P 500 index (in BRL)	35.39%	28.22%
T-bill (BRL)	16.42%	0.10%
		entire sample
unconditional allocation rule, UMV portfolio		
CDI rate		0.11%
IMA-B5		8.48%
IMA-B		38.21%
Ibovespa		5.43%
S&P 500 index (in BRL)		47.12%
T-bill (BRL)		0.65%

Figure A2
Annual wealth index comparison, 2010 to 2024

We show in each panel the wealth index for each calendar year corresponding to investments in the CMV and UMV portfolios, as well as in the CDI rate.

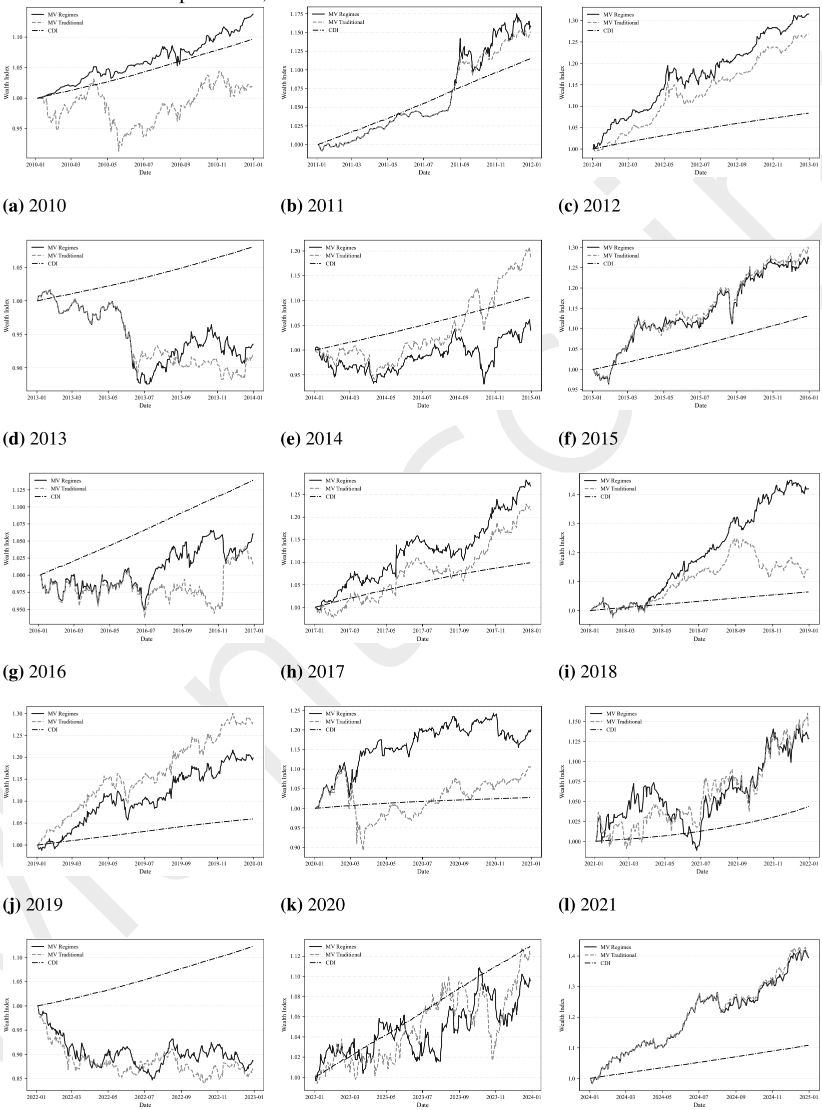

(a) 2010 **(b)** 2011 **(c)** 2012

(d) 2013 **(e)** 2014 **(f)** 2015

(g) 2016 **(h)** 2017 **(i)** 2018

(j) 2019 **(k)** 2020 **(l)** 2021

(m) 2022 **(n)** 2023 **(o)** 2024

Table A3
Robustness checks

The first panel reports the net-of-transaction-costs performance metrics of the CMV portfolio across different regime classifications. In particular, we consider the 50%, 66%, 70% and 80% percentiles of the CDS spread distribution to define the high-risk state. The second and third panels display the differences (Δ) in performance between CMV and UMV portfolios under alternative volatility targets (in percent per annum) and rolling estimation windows (in years).

percentile for regime classification	50%	66%	70%	80%
average return (in % p.a.)	13.29	15.84	13.72	14.53
realized volatility (in % p.a.)	10.65	10.43	10.54	10.39
Sharpe ratio	0.30	0.51	0.33	0.41
Sortino ratio	0.41	0.74	0.48	0.58
maximum drawdown (in %)	-17.60	-16.14	-16.78	-18.75
maximum daily return (in %)	4.70	6.58	6.61	6.61
minimum daily return (in %)	-8.56	-3.76	-4.74	-4.80

volatility target per year	5%	8%	10%	15%	20%
diff in average return (in % p.a.)	0.86%	1.95%	2.08%	2.33%	3.42%
diff in realized volatility (in % p.a.)	-0.43%	-0.51%	-0.43%	-0.31%	-0.09%
diff in Sharpe ratio	0.16	0.22	0.18	0.15	0.18
diff in Sortino ratio	0.25	0.32	0.27	0.21	0.26
diff in maximum drawdown (in %)	2.77%	4.50%	3.15%	-1.29%	-3.40%
diff in turnover (in %)	0.09	0.16	0.19	0.24	0.28

rolling window length (in years)	3	4	5	6	7
diff in average return (in % p.a.)	-2.67%	-0.56%	2.08%	2.01%	-1.01%
diff in realized volatility (in % p.a.)	0.07%	-0.38%	-0.43%	-0.49%	-0.44%
diff in Sharpe ratio	-0.24	-0.04	0.18	0.18	-0.09
diff in Sortino ratio	-0.34	-0.05	0.27	0.25	-0.12
diff in maximum drawdown (in %)	1.19%	10.34%	3.15%	-1.09%	-8.10%
diff in turnover (in %)	0.18	0.18	0.19	0.22	0.32