Large-dimensional portfolio selection in Brazil: an empirical study

Arthur de Souza Calza[a]
Guilherme Valle Moura[a]

[a]Universidade Federal de Santa Catarina (UFSC)

Abstract This article evaluates how alternative models of large-dimensional conditional covariance translate into the performance of out-of-sample portfolios in the Brazilian stock market. Using daily data from B3 for the period 2000–2024 and a rolling window structure, we estimated conditional covariance matrices and mapped them into weights of minimum and mean-variance portfolio strategies, as well as specifications with turnover and positivity constraints on the weights. Performance was evaluated using annualized return, standard deviation, information ratio and turnover. The results highlight the impact of transaction costs on valuation metrics, reinforcing the importance of strategies that favor a combination of stable covariance between the assets with low portfolio turnover.

Keywords: Portfolio selection; large covariance matrices; conditional correlation; stochastic volatility; transaction costs. **JEL codes**: G11, C32, C58.

1. Introduction

Optimal portfolio construction depends critically on the quality of covariance estimates. In minimum and mean–variance settings, optimized weights involve the inverse covariance matrix, so estimation errors in large dimensions can be amplified into extreme and unstable allocations, harming out-of-sample performance and implementability (Markowitz, 1952; Ledoit and Wolf, 2004; Michaud, 1989).

The estimation of high-dimensional covariance matrices is challenging because the number of free parameters grows as $N(N-1)/2$ and the sample covariance matrix becomes noisy and poorly conditioned when N is not small relative to the effective sample size. Regularization methods like shrinkage have been developed to improve conditioning and reduce noise, often delivering superior out-of-sample portfolio performance despite the bias they introduce (Ledoit and Wolf, 2004, 2012, 2017; Fan et al., 2016).

Beyond static regularization, conditional covariance dynamics matter for risk management and portfolio rebalancing. Multivariate conditional volatility

How to cite: Calza, A. S. and Moura, G. V. (). Large-dimensional portfolio selection in Brazil: an empirical study. *Revista Brasileira de Finanças*, , . https://doi.org/
Submitted on March 23, 2026. Published online on .
Editor in charge: .

models such as the dynamic conditional correlation (DCC) framework capture time-varying correlations with a parsimonious parameterization (Engle, 2002; Engle et al., 2019). Matrix-based stochastic volatility formulations, including Wishart stochastic volatility, provide an alternative route that guarantees positive definiteness by construction and can be combined with shrinkage priors to improve stability (Uhlig, 1994, 1997; Kim, 2014).

This paper evaluates these modeling choices in the Brazilian equity market, where concentration, liquidity heterogeneity and regime changes can exacerbate the sensitivity of optimized portfolios to covariance estimation error (Ramalho, 2010; Maciel, 2023). We compare a set of large-dimensional conditional covariance estimators—RiskMetrics, DCC variants with shrinkage, Wishart stochastic volatility (WSV) and its shrunk variant, and a hybrid approximate factor model combined with shrinkage DCC—under portfolio policies of global minimum-variance and mean–variance strategies with a momentum signal.

The empirical design uses daily B3 data (2000–2024) with a rolling window structure. At each rebalancing date, the previous 252 trading days are used to estimate the conditional covariance matrix, which is then mapped into portfolio weights held for the next 21 trading days. Performance is summarized by annualized mean return, volatility, information ratio and turnover, and the analysis incorporates transaction costs, turnover control and long-only constraints.

This paper contributes to the empirical portfolio selection literature by extending the comparison of large-dimensional conditional covariance models to the Brazilian equity market, characterized by limited liquidity and higher volatility. To address these features, the estimation framework relies on shorter rolling windows, highlighting the increased importance of regularization. The analysis also incorporates implementability constraints—transaction costs, turnover penalties, and long-only restrictions—showing that stability is the main driver of out-of-sample performance. In particular, the shrunk-Wishart stochastic volatility (SWSV) model consistently outperforms by avoiding the excessive trading costs associated with high-turnover alternatives.

In addition to this introductory section, this paper is organized as follows. Section 2 provides a brief review of conditional covariance matrix models and the portfolio strategies used in the analysis. Section 3 presents the empirical results obtained with each combination of covariance models and portfolio strategies, evaluating them through metrics such as information ratio and turnover, and verifying the statistical significance of the results. Finally, Section 4 presents the conclusions of this empirical study.

2. Methodology

This section reviews the portfolio policies and large-dimensional conditional covariance estimators used in the empirical analysis.

2.1 Portfolio strategies

2.1.1 Global minimum-variance portfolio

A central benchmark is the *global minimum-variance* (GMV) portfolio, which depends only on the conditional covariance matrix $\mathbf{H}_t$ and therefore isolates the role of covariance modeling (Markowitz, 1952; Ledoit and Wolf, 2004). Let $\mathbf{r}_t$ denote the $N \times 1$ vector of returns and w_t the portfolio weights chosen at $t-1$. The GMV policy solves

$$\min_{w} \quad w'\mathbf{H}_t w \qquad \text{subject to} \qquad w'\mathbf{1} = 1, \tag{1}$$

with closed-form weights

$$w_t^{\mathrm{GMV}} = \frac{\mathbf{H}_t^{-1}\mathbf{1}}{\mathbf{1}'\mathbf{H}_t^{-1}\mathbf{1}}. \tag{2}$$

2.1.2 Mean–variance portfolio with a momentum signal

To introduce a return-seeking component while avoiding full expected-return estimation, we use a mean–variance formulation with an observable momentum signal m_t (Jegadeesh and Titman, 1993; Engle et al., 2019). The portfolio solves

$$\begin{aligned} \min_{w} \quad & w'\mathbf{H}_t w, \\ \text{subject to} \quad & w'\mathbf{1} = 1, \\ & w'm_t = b, \end{aligned} \tag{3}$$

where b is a target exposure to the signal.

The momentum signal m_t is computed as the geometric average of the previous 252 daily returns excluding the most recent 21 returns, i.e., the average over the past year excluding the last month (Engle et al., 2019). In empirical implementations, b can be calibrated from the cross-section of signals; in this paper, we follow Moura et al. (2020) and set b to the average momentum among assets in the top quintile of the momentum ranking.

2.1.3 Turnover control and long-only constraints

A key implementability concern is turnover, since rebalancing costs can overturn frictionless rankings (French, 2008; DeMiguel et al., 2020). Let $turnover_t$ denote one-way turnover at a rebalancing date t, computed from the change in portfolio weights relative to drifted weights w_t^*, as $turnover_t = \sum_{i=1}^{N} |w_{i,t} - w_{i,t}^*|$. Then, we consider a turnover-regularized policy that penalizes deviations from w_t^*:

$$\min_{w} \quad w'\mathbf{H}_t w + \kappa \|w - w_t^*\|_1 \qquad \text{subject to} \qquad w'\mathbf{1} = 1, \tag{4}$$

where $\kappa > 0$ controls the strength of turnover control.

To connect turnover to implementable performance, we also report results under proportional transaction costs. Given a proportional cost rate $c \geqslant 0$, net cumulative return over the holding period is computed as

$$R_{p,t}^{\text{net}} = (1 - c\, turnover_t)(1 + R_{p,t}) - 1, \tag{5}$$

where $R_{p,t}$ is the gross cumulative return (French, 2008; Gârleanu and Pedersen, 2013; Moura et al., 2020).

In addition, we evaluate long-only implementations via the elementwise restriction

$$w_i \geqslant 0, \qquad i = 1, \ldots, N. \tag{6}$$

These constraints act as an implicit regularization device and often reduce the sensitivity of optimized weights to estimation error (Jagannathan and Ma, 2003).

2.2 Conditional covariance matrices

2.2.1 RiskMetrics / exponentially weighted covariance

A widely used baseline is the EWMA covariance estimator popularized by RiskMetrics (J.P. Morgan and Reuters, 1996):

$$\mathbf{H}_t = \lambda \mathbf{H}_{t-1} + (1 - \lambda)\mathbf{r}_{t-1}\mathbf{r}_{t-1}', \qquad \lambda \in (0,1). \tag{7}$$

2.2.2 Dynamic Conditional Correlation (DCC) and shrinkage

DCC models combine univariate volatility dynamics with a parsimonious correlation recursion (Engle, 2002). A standard decomposition is

$$\mathbf{H}_t = \mathbf{D}_t \boldsymbol{\Psi}_t \mathbf{D}_t, \tag{8}$$

where $\mathbf{D}_t$ collects conditional standard deviations and $\mathbf{\Psi}_t$ is the conditional correlation matrix. A convenient parameterization for correlations is

$$\mathbf{\Psi}_t = (1-\eta-\delta)\mathbf{C} + \eta\,\mathbf{s}_{t-1}\mathbf{s}_{t-1}' + \delta\,\mathbf{\Psi}_{t-1}, \tag{9}$$

with standardized returns $\mathbf{s}_t$ and $\eta, \delta \geqslant 0$, $\eta + \delta < 1$.

In large dimensions, the long-run target $\mathbf{C}$ is often regularized through shrinkage. A linear shrinkage form is

$$\widehat{\mathbf{C}} = (1-\rho)\mathbf{S} + \rho\,\mathbf{F}, \qquad \rho \in [0,1], \tag{10}$$

where $\mathbf{S}$ is a sample estimator and $\mathbf{F}$ is a structured target (Ledoit and Wolf, 2004). Nonlinear shrinkage further stabilizes eigenvalues in large-N settings (Ledoit and Wolf, 2012, 2017; Engle et al., 2019).

We consider three DCC specifications that differ in how the long-run correlation target $\mathbf{C}$ is constructed. First, in the baseline DCC model we set $\mathbf{C}$ to the sample correlation matrix computed over the estimation window. Second, in a linear-shrinkage version (DCC-LS) we replace $\mathbf{C}$ by a shrinkage estimator that blends the sample correlation with a structured target. Third, in a nonlinear-shrinkage version (DCC-NLS) we use eigenvalue shrinkage to improve conditioning in large dimensions.

2.2.3 Wishart stochastic volatility and shrinkage

Wishart stochastic-volatility (WSV) models specify stochastic dynamics directly for the conditional covariance matrix $\mathbf{H}_t$ (or, equivalently, for the precision matrix $\mathbf{H}_t^{-1}$), ensuring positive definiteness by construction (Uhlig, 1994, 1997). In the Uhlig-style WSV formulation, the precision matrix evolves through a multiplicative innovation:

$$\mathbf{H}_t^{-1} = \frac{d}{d+1}\,\mathscr{U}(\mathbf{H}_{t-1}^{-1})'\,\mathbf{\Theta}_t\,\mathscr{U}(\mathbf{H}_{t-1}^{-1}), \tag{11}$$

where $d > N-1$ is a degrees-of-freedom parameter. Following Uhlig (1994, 1997), $\mathscr{U}(\mathbf{A})$ denotes the upper-triangular Cholesky factor of a positive definite matrix $\mathbf{A}$ (so that $\mathbf{A} = \mathscr{U}(\mathbf{A})'\,\mathscr{U}(\mathbf{A})$). The innovation matrix $\mathbf{\Theta}_t$ introduces stochastic variation while preserving symmetry and positive definiteness.

In practice, WSV is often implemented through its analytic filtering recursion, which can be expressed as an EWMA-type update for a scale matrix $\mathbf{S}_t$:

$$\mathbf{S}_t = \frac{d}{d+1}\mathbf{S}_{t-1} + \frac{1}{d+1}\mathbf{r}_t\mathbf{r}_t', \tag{12}$$

with $d > N-1$. Under the WSV filtering representation, the scale matrix admits a direct interpretation as the filtered mean covariance, $\mathbb{E}[\mathbf{H}_t] = \mathbf{S}_t$. This makes the role of d transparent: defining $\lambda = d/(d+1)$ yields $\mathbf{S}_t = \lambda\,\mathbf{S}_{t-1} + (1-\lambda)\,\mathbf{r}_t\mathbf{r}_t'$, so larger d implies λ closer to one and therefore more persistence and smoother covariance dynamics.

In high dimensions, portfolio policies depend on $\mathbf{H}_t^{-1}$ and therefore require numerically stable inversion. We consider a shrunk-Wishart stochastic-volatility variant (SWSV) that regularizes the filter output by shrinking the covariance scale toward a diagonal (or otherwise structured) target before inversion, reducing the impact of noisy sample correlations and improving conditioning. For further details on the WSV filtering representation and related shrinkage variants, see Uhlig (1997).

2.2.4 Approximate factor models

Approximate factor models reduce effective dimension by decomposing returns into a low-rank factor component and an idiosyncratic component. Under a single-factor specification,

$$\mathbf{r}_t = \mathbf{A} + \mathbf{B} f_t + \mathbf{u}_t, \tag{13}$$

implying the conditional covariance decomposition

$$\mathbf{H}_t = \sigma_f^2\,\mathbf{B}\mathbf{B}' + \mathbf{\Sigma}_{u,t}. \tag{14}$$

Following De Nard et al. (2020), we consider a hybrid AFM–DCC–NLS approach that models the idiosyncratic covariance $\mathbf{\Sigma}_{u,t}$ using shrinkage DCC techniques (Engle et al., 2019; Ledoit and Wolf, 2017).

3. Results

This section reports the empirical performance of the portfolio policies and conditional covariance estimators reviewed in Section 2, focusing on their out-of-sample implications for portfolio selection in a large-dimensional setting. The discussion follows the empirical logic emphasized in Moura et al. (2020). Since optimized weights depend on the inverse conditional covariance matrix, differences in covariance dynamics feed directly into realized volatility, turnover, and risk-adjusted performance.

This empirical analysis is based on daily historical trading data for equities listed on the Brazilian stock exchange B3, covering the period from 2000 to 2024. To evaluate high-dimensional portfolio performance, we selected $N = 100$ assets with complete historical data, and the a rolling-window design.

At each portfolio formation date, the previous 252 trading days (approximately one year) are used to estimate the conditional covariance matrix, which then serves as an input to the portfolio selection strategies. The resulting portfolio is held (and evaluated) over the next 21 trading days (approximately one month), and the evaluation metrics are computed over that out-of-sample horizon. The choice of a one-year estimation window and a one-month evaluation window is similar to the empirical study by Moura et al. (2020), but with reductions in the estimation window.

This reduction, both in the portfolio size $N = 100$ and in the window $T = 252$, proved necessary given the specific conditions of the Brazilian capital market. Moura et al. (2020) used historical data from the US stock exchange to develop portfolios with $N = 100$, 500, 1000 and estimation windows with $T = 1250$. In comparison, a window with $T = 252$ rarely formed a portfolio with $N = 100$ and with a complete history during the window (a necessary condition for estimating a stable conditional covariance matrix), in data prior to 2005. Only after this period was it possible to develop portfolios with a complete history and $N = 100$, but with a reduced estimation window.

Portfolio performance was evaluated using annualized returns, annualized standard deviations, an information ratio, and turnover. Returns were computed out of sample over each 21-day evaluation window. For comparability across strategies, cumulative returns and standard deviations were reported in annualized terms, using the same scaling for all models. Volatility is central not only as a direct measure of risk, but also because optimization depends on the estimated precision matrix, so estimation noise can be amplified into unstable weights and unstable realized risk (Michaud, 1989; Ledoit and Wolf, 2004). In high-dimensional settings, volatility comparisons should be interpreted jointly with the stability of allocations. Even if two estimators deliver similar average variances, the one that induces smoother weight paths typically entails lower trading needs and therefore better implementability. This stability channel is emphasized in Moura et al. (2020) and is consistent with the view that improved conditioning and regularization of covariance estimates reduce extreme and unstable portfolio positions (Ledoit and Wolf, 2004; Engle et al., 2019).

Risk-adjusted performance was summarized through an information ratio computed from the out-of-sample return series. Let $r_{p,t}$ denote the portfolio return observation used for evaluation (under the rolling-window design), and let μ_p and σ_p denote its mean and standard deviation over the out-of-sample sample. The information ratio is defined as $IR_p = \mu_p/\sigma_p$, so it summarizes realized return per unit of realized volatility and remains directly tied to out-of-sample outcomes (Engle et al., 2019; Moura et al., 2020).

The metrics should be interpreted jointly rather than in isolation. For minimum-variance portfolios, lower volatility is a direct measure of success, while turnover and the information ratio complement volatility by capturing implementability and return per unit of realized risk.

The tables below report point estimates for realized volatility and information ratios. Following Moura et al. (2020) and the robust testing framework in Ledoit and Wolf (2008) and Ledoit and Wolf (2011), statistical inference is conducted using influence-function-based tests with prewhitened HAC standard errors. Within each portfolio policy and return specification, models are compared to a benchmark defined by the best point estimate: the smallest sample standard deviation for volatility and the largest sample information ratio for risk-adjusted performance. The hypotheses are one-sided and aligned with this benchmarking logic: volatility is tested for being larger than the minimum-volatility benchmark and information ratios are tested for being smaller than the maximum-IR benchmark.

Tables 1–16 summarize statistical significance with asterisks. One, two, and three asterisks correspond to rejection at the 90%, 95%, and 99% confidence levels, respectively. In the volatility columns, asterisks flag models whose standard deviation is significantly higher than the minimum-volatility benchmark in the corresponding table; in the information ratio columns, they flag models whose information ratio is significantly lower than the maximum-IR benchmark. The benchmark itself, and any models that are not statistically different from it, are highlighted in bold.

3.1 Minimum-variance portfolios

The minimum-variance results are reported for two implementations, the baseline GMV portfolio and a turnover-controlled variant that penalizes large reallocations between rebalancing dates. Since both portfolios share the same objective (risk minimization), comparisons across conditional covariance estimators focus on realized volatility and turnover.

The results in Tables 1 and 2 highlight a pattern also documented in the high-dimensional portfolio literature: because minimum-variance weights depend on the inverse covariance matrix, estimation noise in the covariance input can be amplified by the optimizer and manifest as unstable allocations (Michaud, 1989; Ledoit and Wolf, 2004; Engle et al., 2019; Moura et al., 2020). In this context, lower turnover indicates that the covariance estimate, and therefore its inverse, evolves smoothly enough to avoid frequent portfolio reorientation.

Table 1
Results (no transaction costs): GMV portfolio.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	**0.13**	45.22	102.64***	1.30
SWSV	0.05*	4.25	**25.72**	0.08
DCC	0.03**	3.10	**34.96**	0.17
DCC-LS	0.01**	1.19	**33.45**	0.14
DCC-NLS	-0.00**	-0.00	**27.66**	0.10
AFM-DCC-NLS	-0.03***	-2.80	**28.41**	0.13
WSV	-0.03**	-3.16	30.54***	0.13

Table 2
Results (no transaction costs): minimum-variance portfolio with turnover control.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	**0.12**	28.95	71.52***	1.03
SWSV	**0.05**	4.23	**25.71**	0.08
DCC	0.03*	3.07	**34.87**	0.16
DCC-LS	0.01**	1.16	**33.40**	0.14
DCC-NLS	-0.00**	-0.02	**27.64**	0.10
AFM-DCC-NLS	-0.03**	-2.80	**28.36**	0.13
WSV	-0.03**	-3.18	30.52***	0.13

Across both implementations, the most stable results emerged under the SWSV specification. This estimator produced the lowest annualized volatility (approximately 25.7%) and the lowest turnover (approximately 0.08) in the minimum-variance strategies. The dominance is remarkably consistent across the baseline and turnover-controlled implementations, suggesting that the main driver could be the covariance input itself. When the covariance estimate changes gradually, the risk-minimizing allocation also changes gradually, mechanically reducing trading needs.

Conversely, RiskMetrics is associated with extremely high turnover and volatility. Even though the portfolio's objective is risk minimization, the RiskMetrics input produced annualized volatility above 70% in the implementation with turnover control and above 100% in the base version, with turnover around 1.03 to 1.30. A fast-reacting covariance estimate leads the optimizer to chase short-duration fluctuations in joint movements, producing unstable weights and high trading intensity. Empirical volatility patterns are consistent with a specification that relies primarily on recent past returns to update the conditional covariance structure (Michaud, 1989; Moura et al., 2020).

The DCC family lies between these extremes. When the long-term correlation target was estimated with the sample target, both volatility and turnover were comparatively higher than the shrinkage DCC portfolios. Introducing shrinkage improved stability, which is consistent with the role of shrinkage in improving the conditioning of large-dimension covariance estimates (Ledoit and Wolf, 2004, 2017; Engle et al., 2019). In particular, the nonlinear-shrinkage variant delivered materially lower volatility and turnover than the sample-target version.

To compare models within the minimum-variance rule, it is informative to focus on risk-adjusted performance and implementability. On this dimension, the information ratio tends to favor specifications that either deliver lower realized volatility for comparable holding-period returns or, equivalently, avoid excessive risk-taking induced by unstable covariance inverses. The SWSV specification achieved a positive information ratio with the lowest realized volatility and turnover. This result was also found, for the US market, in the study by Moura et al. (2020), indicating that the combination of SWSV and minimum variance strategies is a useful specification when the goal is to minimize turnover.

Table 3
Transaction-cost sensitivity: global minimum-variance policy.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-11.05	116.26***	**-0.03**
SWSV	-0.57	**24.43**	**-0.01**
DCC	-7.04	32.77*	**-0.06**
DCC-LS	-7.21	**31.17**	**-0.07**
DCC-NLS	-6.35	**26.54**	-0.07*
AFM-DCC-NLS	-10.47	27.80*	-0.11***
WSV	-10.67	29.63***	-0.10**

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-67.32	175.79***	**-0.11**
SWSV	-5.38	**23.35**	**-0.07**
DCC	-17.18	31.83**	-0.16*
DCC-LS	-15.60	29.78*	-0.15*
DCC-NLS	-12.70	**25.79**	-0.14**
AFM-DCC-NLS	-18.14	27.63*	-0.19***
WSV	-18.18	29.33***	-0.18***

Table 3 and Table 4 show how quickly performance deteriorates for high-turnover specifications as costs rise. In these cases, low-turnover inputs remain comparatively resilient, while RiskMetrics becomes unattractive once trading frictions are introduced. Alternatively, SWSV becomes the most attractive option also in terms of variance, performing with the lowest variance in both the unrestricted and turnover-restricted strategies.

Table 4
Transaction-cost sensitivity: minimum-variance policy with turnover control.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-16.76	104.96***	**-0.05**
SWSV	-0.58	**24.42**	**-0.01**
DCC	-6.98	32.68*	**-0.06**
DCC-LS	-7.18	**31.13**	**-0.07**
DCC-NLS	-6.35	**26.52**	-0.07*
AFM-DCC-NLS	-10.44	27.75*	-0.11***
WSV	-10.68	29.63***	-0.10**

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-62.47	172.48***	**-0.10**
SWSV	-5.39	**23.34**	**-0.07**
DCC	-17.03	31.71**	-0.16*
DCC-LS	-15.53	29.73*	-0.15*
DCC-NLS	-12.68	**25.77**	-0.14**
AFM-DCC-NLS	-18.08	27.57*	-0.19***
WSV	-18.19	29.33***	-0.18***

The same minimum-variance strategies were also evaluated under a long-only constraint, which restricted weights to be nonnegative. This restriction reduced turnover in all model–strategy combinations and, in most cases, also reduced realized volatility relative to Tables 1 and 2. In the baseline GMV policy, turnover moved from the order of 0.08–1.30 to approximately 0.03–0.04 across all covariance inputs. This reduction was particularly pronounced under RiskMetrics, whose turnover fell from 1.30 to 0.04 and whose annualized volatility dropped from 102.64% to 38.52%.

In terms of risk, the long-only minimum-variance results displayed a clear improvement in dispersion for most conditional covariance models. Within the DCC family, annualized standard deviation declined materially relative to the unconstrained case, and the sample-target DCC delivered the lowest volatility among long-only minimum-variance portfolios when transaction costs were ignored. In Table 5, DCC achieved 25.40% annualized volatility, slightly below SWSV at 26.34%. The two notable exceptions to the variance-reduction pattern were SWSV and AFM-DCC-NLS. Both specifications displayed a small increase in volatility under long-only restrictions, with SWSV moving from 25.72% to 26.34% and AFM-DCC-NLS moving from 28.41% to 30.20%. These deviations suggested that, for these covariance dynamics, the unconstrained optimizer had already produced stable allocations, and the nonnegativity constraint limited risk-reducing long–short adjustments.

Tables 5 and 6 also indicate that the additional turnover penalty became less relevant once long-only constraints were imposed. Since turnover is already clustered around 0.03–0.04, the turnover-controlled objective produced only marginal additional changes in realized outcomes.

Table 5
Results (no transaction costs): GMV portfolio, long-only.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	**0.03**	3.60	38.52**	0.04
SWSV	**0.08**	7.00	**26.34**	0.04
DCC	-0.02***	-1.64	**25.40**	0.03
DCC-LS	-0.01**	-0.93	**26.63**	0.03
DCC-NLS	-0.01**	-0.95	27.37*	0.03
AFM-DCC-NLS	**0.05**	5.54	30.20**	0.04
WSV	**0.06**	6.03	**29.32**	0.04

Table 6
Results (no transaction costs): minimum-variance portfolio with turnover control, long-only.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	**0.03**	3.60	38.50**	0.04
SWSV	**0.08**	7.00	**26.33**	0.04
DCC	-0.02***	-1.65	**25.39**	0.03
DCC-LS	-0.01**	-0.94	**26.61**	0.03
DCC-NLS	-0.01**	-0.95	27.36*	0.03
AFM-DCC-NLS	**0.05**	5.53	30.19**	0.04
WSV	**0.06**	6.03	**29.32**	0.04

Tables 7 and 8 extend the transaction-cost exercise to long-only portfolios. The impact of transaction costs is materially smaller than in Tables 3 and 4, which is consistent with the large turnover reductions induced by nonnegative weights. Under RiskMetrics and $c = 0.10$, net return improved from -67.32% in the unconstrained GMV policy to -2.04% under long-only. Under SWSV, net return at $c = 0.10$ reached 1.55%, while WSV delivered 1.25%.

Table 7
Transaction-cost sensitivity: global minimum-variance policy, long-only.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	0.78	36.16*	**0.01**
SWSV	4.28	**25.16**	**0.05**
DCC	-3.70	**24.91**	-0.04***
DCC-LS	-3.07	**25.91**	-0.03**
DCC-NLS	-3.11	26.56*	-0.03**
AFM-DCC-NLS	2.88	28.89*	**0.03**
WSV	3.64	**27.84**	**0.04**

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-2.04	**33.90**	**-0.02**
SWSV	1.55	**24.07**	**0.02**
DCC	-5.76	**24.48**	-0.07**
DCC-LS	-5.21	**25.27**	-0.06**
DCC-NLS	-5.27	**25.82**	-0.06**
AFM-DCC-NLS	0.22	27.68*	**0.00**
WSV	1.25	**26.45**	**0.01**

Table 8
Transaction-cost sensitivity: minimum-variance policy with turnover control, long-only.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	0.78	36.14*	**0.01**
SWSV	4.28	**25.16**	**0.05**
DCC	-3.71	**24.90**	-0.04***
DCC-LS	-3.08	**25.89**	-0.03**
DCC-NLS	-3.11	26.55*	-0.03**
AFM-DCC-NLS	2.87	28.89*	**0.03**
WSV	3.64	**27.84**	**0.04**

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-2.04	**33.87**	**-0.02**
SWSV	1.55	**24.07**	**0.02**
DCC	-5.77	**24.47**	-0.07**
DCC-LS	-5.22	**25.25**	-0.06**
DCC-NLS	-5.27	**25.82**	-0.06**
AFM-DCC-NLS	0.21	27.68*	**0.00**
WSV	1.25	**26.45**	**0.01**

3.2 Mean–variance with momentum portfolios

Relative to the minimum-variance rule, the mean-variance portfolios with a momentum signal depend not only on the conditional covariance matrix but also on an expected-return proxy, which changes the empirical trade-offs across models. In this setting, the conditional covariance estimator affects the risk component of the objective, while the momentum signal introduces an additional source of time variation in optimal weights.

Table 9
Results (no transaction costs): mean–variance portfolio with momentum.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	**0.07**	22.51	87.36***	0.83
SWSV	**0.01**	0.53	**26.47**	0.12
DCC	-0.03*	-2.92	30.68**	0.22
DCC-LS	-0.05**	-4.73	**28.24**	0.17
DCC-NLS	-0.06**	-5.17	**26.03**	0.13
AFM-DCC-NLS	-0.04**	-4.50	30.15**	0.16
WSV	-0.05**	-5.22	32.53*	0.16

Table 10
Results (no transaction costs): mean–variance portfolio with momentum and turnover control.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	**0.06**	15.00	69.67***	0.64
SWSV	**0.01**	0.49	**26.45**	0.12
DCC	-0.03*	-2.88	30.63*	0.22
DCC-LS	-0.05**	-4.76	**28.19**	0.17
DCC-NLS	-0.06**	-5.21	**26.02**	0.13
AFM-DCC-NLS	-0.04*	-4.54	30.14**	0.16
WSV	-0.05*	-5.25	32.51*	0.16

Tables 9 and 10 summarize the results for the baseline mean–variance implementation and for the turnover-controlled variant, without taking transaction costs into account. RiskMetrics delivered the highest information ratio in both cases, with $IR = 0.07$ in the baseline portfolio and $IR = 0.06$ under turnover control. The same specification also produced the highest volatility and the highest turnover among all covariance inputs. In the baseline portfolio, annualized volatility reached 87.36% with turnover of 0.83, while the turnover-controlled variant reduced these figures to 69.67% and 0.64. These outcomes were consistent with the idea that a fast-moving covariance input can generate aggressive changes in weights when combined with a return signal, improving realized risk-adjusted performance in frictionless terms but at the cost of substantially higher dispersion and rebalancing needs.

The DCC family exhibited intermediate behavior, with the expected stability gains from shrinkage showing up primarily through lower turnover and lower dispersion. The sample-target DCC produced turnover of 0.22 and annualized volatility around 30.7%, with an information ratio of -0.03. Introducing shrinkage improved stability: DCC-LS reduced annualized standard deviation to 28.24% and turnover to 0.17, while DCC-NLS further reduced volatility to 26.03% and turnover to 0.13. Despite these stability improvements, all DCC variants delivered negative information ratios in the sample, with annualized returns ranged from -2.92% for the sample-target DCC to -5.17% for DCC-NLS. These strategies performed below those obtained with the minimum variance strategy, indicating that the return component introduced through momentum did not offset the realized risk in these specifications. The AFM-DCC-NLS extension also did not improve performance in this setting. It delivered an information ratio of -0.04 with annualized return of -4.50%, volatility of 30.15%, and turnover of 0.16, which placed it closer to the higher-dispersion configurations rather than to the most stable shrinkage outcome. Under turnover control, AFM-DCC-NLS remained essentially unchanged, with $IR = -0.04$, mean return of -4.54%, annualized standard deviation of 30.14%, and turnover of 0.16.

Within the Wishart stochastic-volatility specifications, SWSV produced the smoothest allocations and the lowest dispersion, despite a much lower information ratio than RiskMetrics. In both implementations, SWSV delivered annualized standard deviation close to 26.5% and turnover around 0.12, while annualized returns stayed near 0.5% and the information ratio remained near 0.01. The non-shrunk Wishart stochastic-volatility model, WSV, performed materially worse than its shrunk counterpart. It delivered negative information ratios around -0.05, with annualized returns near -5.2%, volatility around 32.5%, and turnover around 0.16. Under turnover control, WSV also underperformed its shrinkage version, with $IR = -0.05$, mean return of -5.25%, annualized standard deviation of 32.51%, and turnover of 0.16. This gap between SWSV and WSV highlighted the empirical role of shrinkage within Wishart-based dynamics. Smoother covariance paths translated into smoother allocations and lower dispersion, while the unshrunk specification induced higher realized volatility and higher trading intensity.

Once transaction costs were introduced, the ranking changed sharply because net performance became sensitive to turnover. In the baseline portfolio, net annualized return under RiskMetrics dropped to -20.20% for $c = 0.05$ and to -62.91% for $c = 0.10$, while SWSV yielded -6.44% and -13.40% over the same cost levels. This ranking reversal mirrored the empirical mechanism already observed in minimum-variance portfolios. High-turnover covariance inputs could look attractive without frictions, but they became dominated once trading costs were priced in.

Table 11
Transaction-cost sensitivity: mean–variance portfolio with momentum.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-20.20	66.81***	**-0.09**
SWSV	-6.44	**26.19**	**-0.07**
DCC	-15.31	32.20*	**-0.14**
DCC-LS	-14.60	**29.38**	**-0.14**
DCC-NLS	-12.96	**26.58**	-0.14**
AFM-DCC-NLS	-13.97	30.24**	-0.13**
WSV	-14.46	32.97***	-0.13***

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-62.91	57.79***	-0.31***
SWSV	-13.40	**26.26**	**-0.15**
DCC	-27.71	35.08**	-0.23*
DCC-LS	-24.46	31.23*	-0.23*
DCC-NLS	-20.75	**27.44**	-0.22**
AFM-DCC-NLS	-23.43	30.97**	-0.22**
WSV	-23.70	34.09***	-0.20***

Table 11 and Table 12 report the same sensitivity exercise for the baseline and turnover-controlled implementations. Turnover control partially mitigated the loss in net performance, but it did not overturn the main comparative pattern. Under turnover control, RiskMetrics improved from -62.91% to -52.00% at $c = 0.10$, yet it still lagged far behind lower-turnover alternatives such as SWSV at -13.42%. Across DCC specifications, the same ordering persisted under costs, and DCC-NLS remained more resilient than DCC-LS and the sample-target DCC because it induced lower turnover and therefore smaller net-return penalties.

The mean–variance strategies were also re-estimated under a long-only constraint. In this case, turnover declined in all combinations of covariance inputs and portfolio rules, and information ratios improved for almost all conditional covariance models relative to Tables 9 and 10. The notable exception was RiskMetrics. Without long-only restrictions it delivered the highest

Table 12
Transaction-cost sensitivity: mean–variance portfolio with momentum and turnover control.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-18.50	57.25***	**-0.09**
SWSV	-6.46	**26.18**	**-0.07**
DCC	-15.17	32.10*	**-0.14**
DCC-LS	-14.58	**29.32**	**-0.14**
DCC-NLS	-12.98	**26.56**	-0.14**
AFM-DCC-NLS	-13.99	30.23**	-0.13**
WSV	-14.49	32.96***	-0.13***

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-52.00	52.04***	-0.29**
SWSV	-13.42	**26.26**	**-0.15**
DCC	-27.47	34.90**	-0.23*
DCC-LS	-24.40	31.16*	-0.23*
DCC-NLS	-20.74	**27.42**	-0.22**
AFM-DCC-NLS	-23.44	30.95**	-0.22**
WSV	-23.72	34.09***	-0.20***

information ratio, but under long-only it became the weakest specification. In Table 13, volatility fell sharply from 87.36% to 24.97%, but mean return moved from 22.51% to -2.59%, which drove the information ratio from 0.07 to -0.03.

Among the remaining models, SWSV delivered the highest information ratio in the long-only mean–variance portfolios. It achieved $IR = 0.04$ with annualized return of 3.17% and volatility of 25.82%, together with turnover close to 0.05. WSV also improved markedly relative to its unconstrained counterpart. Its information ratio moved from -0.05 to 0.03, with annualized return of 2.48% and volatility of 27.91%. Within the DCC family, information ratios turned closer to zero or mildly positive under long-only mainly through lower turnover, while annualized volatility remained near 28%.

Tables 15 and 16 report transaction-cost sensitivity under long-only restrictions. The reduction in turnover mitigated the net-return deterioration relative to Tables 11 and 12. Under RiskMetrics, net return at $c = 0.10$ improved from -62.91% in the unconstrained portfolio to -7.78% under long-only, but performance remained weak because gross returns were already negative. At $c = 0.05$, SWSV delivered net return of 0.48% with $IR = 0.01$, while WSV remained close to zero. At $c = 0.10$, net returns turned negative for all models,

Table 13
Results (no transaction costs): mean–variance portfolio with momentum, long-only.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	-0.03*	-2.59	**24.97**	0.04
SWSV	**0.04**	3.17	**25.82**	0.05
DCC	**0.01**	0.92	28.41**	0.04
DCC-LS	**0.00**	0.45	28.21**	0.04
DCC-NLS	-0.01*	-0.66	28.36**	0.04
AFM-DCC-NLS	**0.03**	2.62	30.11***	0.04
WSV	**0.03**	2.48	27.91*	0.04

Table 14
Results (no transaction costs): mean–variance portfolio with momentum and turnover control, long-only.

Covariance model	IR	Mean return (%)	Std. dev. (%)	Turnover
RiskMetrics	-0.03*	-2.57	**24.97**	0.04
SWSV	**0.04**	3.17	**25.81**	0.05
DCC	**0.01**	0.92	28.41**	0.04
DCC-LS	**0.00**	0.46	28.21**	0.04
DCC-NLS	-0.01*	-0.66	28.36**	0.04
AFM-DCC-NLS	**0.03**	2.62	30.10***	0.04
WSV	**0.03**	2.48	27.91*	0.04

but the magnitudes remained substantially smaller than in the unconstrained case.

Table 15
Transaction-cost sensitivity: mean–variance portfolio with momentum, long-only.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-5.18	**24.89**	-0.06*
SWSV	0.48	**25.80**	**0.01**
DCC	-1.57	28.37**	**-0.02**
DCC-LS	-2.04	28.18**	**-0.02**
DCC-NLS	-3.15	28.33**	**-0.03**
AFM-DCC-NLS	0.12	30.05***	**0.00**
WSV	-0.02	27.96*	**-0.00**

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-7.78	**24.83**	-0.09**
SWSV	-2.20	**25.80**	**-0.02**
DCC	-4.05	28.34**	**-0.04**
DCC-LS	-4.53	28.16**	**-0.05**
DCC-NLS	-5.64	28.32**	-0.06*
AFM-DCC-NLS	-2.39	30.00***	**-0.02**
WSV	-2.53	28.01*	**-0.03**

Table 16
Transaction-cost sensitivity: mean–variance portfolio with momentum and turnover control, long-only.

$c = 0.05$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-5.17	**24.89**	-0.06*
SWSV	0.49	**25.80**	**0.01**
DCC	-1.56	28.37**	**-0.02**
DCC-LS	-2.04	28.18**	**-0.02**
DCC-NLS	-3.15	28.33**	**-0.03**
AFM-DCC-NLS	0.12	30.05***	**0.00**
WSV	-0.02	27.95*	**-0.00**

$c = 0.10$			
Covariance model	Net return (%)	Std. dev. (%)	IR
RiskMetrics	-7.76	**24.83**	-0.09**
SWSV	-2.20	**25.80**	**-0.02**
DCC	-4.05	28.34**	**-0.04**
DCC-LS	-4.53	28.16**	**-0.05**
DCC-NLS	-5.64	28.32**	-0.06*
AFM-DCC-NLS	-2.39	30.00***	**-0.02**
WSV	-2.53	28.01*	**-0.03**

4. Conclusion

This paper evaluated how alternative large-dimensional conditional covariance models translate into out-of-sample portfolio performance in the Brazilian equity market. Using daily B3 data and a rolling-window design, conditional covariance matrices were estimated over a fixed historical window and mapped into portfolio weights. These weights were held over a short evaluation horizon, and performance was evaluated based on annualized mean return, volatility, information ratio, and turnover. The analysis compared standard benchmarks (RiskMetrics and DCC-type specifications) with Wishart stochastic-volatility models, with special attention to the SWSV shrunk-Wishart specification, and assessed how rankings changed once proportional transaction costs and practical constraints were introduced.

For global minimum-variance portfolios, the main empirical message was that SWSV consistently delivered the most implementable risk control. In the frictionless evaluation, SWSV was among the lowest-volatility allocations and simultaneously among the lowest-turnover allocations, a combination that is precisely the stability channel emphasized in the recent literature on large-dimensional covariance modeling and portfolio choice. This pattern closely matched the evidence reported in Moura et al. (2020): smoother correlation dynamics translated into smoother optimal weights and lower trading needs, without sacrificing risk control. When transaction costs were considered, the advantage of low turnover became decisive for net performance, and SWSV remained a robust choice because it avoided the large net-return penalties induced by high-turnover covariance inputs. This stability is especially valuable in emerging environments like the Brazilian capital market, where asset concentration and liquidity gaps can otherwise magnify estimation noise into extreme rebalancing costs.

For mean-variance portfolios with a momentum signal, the results reinforced that the covariance input matters not only through risk, but also through how it interacts with a time-varying return proxy. In frictionless terms, some specifications could deliver higher information ratios while also generating very high volatility and high turnover, making their performance fragile from an implementation perspective. In contrast, SWSV again stood out for producing comparatively stable allocations, with low dispersion and low turnover, and it became particularly competitive once transaction costs were introduced. Overall, the evidence indicated that, in this setting, the return signal did not eliminate the need for regularized covariance dynamics; rather, it made stability even more important because mean-variance weights can react strongly to changes in inputs.

Introducing an explicit turnover penalty produced improvements that were directionally consistent across both portfolio problems. Turnover control reduced trading intensity and mitigated the deterioration of net returns under costs, especially for covariance estimators that otherwise induced frequent reallocations. At the same time, the turnover-penalized objective did not overturn the main comparative message: models that were already stable in their implied covariance dynamics (most notably SWSV) remained among the best-performing choices once implementability was accounted for, while high-turnover specifications remained vulnerable to transaction costs even after partial stabilization.

Long-only constraints provided an additional and practically relevant form of regularization. Imposing nonnegative weights sharply reduced turnover across model-strategy combinations and generally improved the realism of the optimized portfolios by limiting extreme long-short adjustments. Under long-only, the incremental gains from turnover control became smaller, indicating that the constraint itself absorbed part of the stabilization role that the turnover penalty was designed to provide. Importantly, long-only restrictions also reduced the sensitivity of net performance to transaction costs, strengthening the general conclusion that portfolio evaluation in high dimensions should prioritize implementability metrics jointly with volatility and risk-adjusted performance.

Future work can extend the analysis in directions that are directly aligned with implementation constraints faced by portfolio managers. A natural next step is to study portfolio policies under additional feasibility restrictions, such as leverage limits and related exposure constraints, which are central in practice and can materially change the mapping from covariance estimates to weights. In addition, testing further covariance specifications and regularization schemes–including alternative shrinkage targets, factor structures, and hybrid dynamics–could provide a more complete view of how the quality of risk-matrix estimation affects realized performance in the Brazilian market, particularly when the objective is robust, low-turnover portfolio construction.

Conflict of interest The authors declare no conflict of interest.

Data availability The data and code are available from the authors. For those interested, please contact us via email: arthurcalza@hotmail.com.

References

De Nard, G., Ledoit, O. and Wolf, M. (2020). Factor models for portfolio selection in large dimensions: The good, the better and the ugly, *Journal of Financial Econometrics*, . Forthcoming.

DeMiguel, V., Martín-Utrera, A., Nogales, F. J. and Uppal, R. (2020). A transaction-cost perspective on the multitude of firm characteristics, *The Review of Financial Studies*, *33*(5), 2180–2222.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models, *Journal of Business and Economic Statistics*, *20*, 339–350.

Engle, R. F., Ledoit, O. and Wolf, M. (2019). Large dynamic covariance matrices, *Journal of Business and Economic Statistics*, *37*, 363–375.

Fan, J., Liao, Y. and Liu, H. (2016). An overview of the estimation of large covariance and precision matrices, *The Econometrics Journal*, *19*(1), C1–C32.

French, K. R. (2008). Presidential address: The cost of active investing, *Journal of Finance*, *63*(4), 1537–1573.

Gârleanu, N. and Pedersen, L. H. (2013). Dynamic trading with predictable returns and transaction costs, *Journal of Finance*, *68*(6), 2309–2340.

Jagannathan, R. and Ma, T. (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps, *Journal of Finance*, *58*(4), 1651–1684.

Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance*, *48*, 65–91.

J.P. Morgan and Reuters (1996). *RiskMetrics Technical Document*.

Kim, D. (2014). Maximum likelihood estimation for vector autoregressions with multivariate stochastic volatility, *Economics Letters*, *123*(2), 282–286.

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices, *Journal of Multivariate Analysis*, *88*(2), 365–411.

Ledoit, O. and Wolf, M. (2008). Robust performance hypothesis testing with the sharpe ratio, *Journal of Empirical Finance*, *15*, 850–859.

Ledoit, O. and Wolf, M. (2011). Robust performance hypothesis testing with the variance, *Working paper*, University of Zurich.

Ledoit, O. and Wolf, M. (2012). Nonlinear shrinkage estimation of large-dimensional covariance matrices, *The Annals of Statistics*, *40*, 1024–1060.

Ledoit, O. and Wolf, M. (2017). Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets Goldilocks, *The Review of Financial Studies*, *30*(12), 4349–4388.

Maciel, L. d. S. (2023). Brazilian stock-market efficiency before and after covid-19: The roles of fractality and predictability, *Global Finance Journal*, *58*, 100887.
URL: https://doi.org/10.1016/j.gfj.2023.100887

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

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

Moura, G. V., Santos, A. A. P. and Ruiz, E. (2020). Comparing high dimensional conditional covariance matrices: Implications for portfolio selection, *Journal of Banking and Finance*, *118*.

Ramalho, C. (2010). Toward convergence? the brazilian stock market concentration and the role of private equity in its diversification. First draft (July 2010). Preliminary draft; please do not cite without permission.
URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1961642

Uhlig, H. (1994). On singular wishart and singular multivariate beta distributions, *The Annals of Statistics*, *22*, 395–405.

Uhlig, H. (1997). Bayesian vector autoregressions with stochastic volatility, *Econometrica*, *65*, 59–74.