UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO EM COMPUTAÇÃO

KEVIN GABRIEL RAMISCH PERGHER

Orthogonal Hierarchical Risk Parity Methods for Portfolio Allocation in the Brazilian Stock Market

Thesis presented in partial fulfillment of the requirements for the degree of Master of Computer Science

Advisor: Prof. Dr. Jacob Scharcanski

Porto Alegre
August 2026

CIP — CATALOGING-IN-PUBLICATION

Pergher, Kevin Gabriel Ramisch

Orthogonal Hierarchical Risk Parity Methods for Portfolio Allocation in the Brazilian Stock Market / Kevin Gabriel Ramisch Pergher. – Porto Alegre: PPGC da UFRGS, 2026.

89 f.: il.

Thesis (Master) – Universidade Federal do Rio Grande do Sul. Programa de Pós-Graduação em Computação, Porto Alegre, BR–RS, 2026. Advisor: Jacob Scharcanski.

1. Portfolio allocation. 2. Hierarchical risk parity. 3. Orthogonal projections. 4. Brazilian stock market. 5. Machine learning. I. Scharcanski, Jacob. II. Título.

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
Reitora: Prof[a]. Márcia Cristina Bernardes Barbosa
Vice-Reitor: Prof. Pedro de Almeida Costa
Pró-Reitora de Pós-Graduação: Prof[a]. Cláudia Wasserman
Diretor do Instituto de Informática: Prof. Luciano Paschoal Gaspary
Coordenador do PPGC: Prof. Gabriel Luca Nazar
Bibliotecária-chefe do Instituto de Informática: Alexsander Borges Ribeiro

“The information you have is not the information you want.
The information you want is not the information you need.
The information you need is not available.
The information available costs more than you want to pay.”

— Peter L. Bernstein

AGRADECIMENTOS

Meus agradecimentos iniciais ao corpo técnico e docente do Instituto de Informática da UFRGS, em especial ao professor Jacob Scharcanski, que muito auxiliaram, apoiaram e acreditaram no desenvolvimento desta dissertação, viabilizando minha participação em congressos nacionais e internacionais e as publicações relacionadas. Estendo esse agradecimento a muitos professores e conhecidos da Faculdade de Economia e do Instituto de Física, os quais muito colaboraram na discussão e na formulação das mais diversas ideias no decorrer da minha trajetória.

Agradeço fortemente o apoio incondicional da minha família, com muito carinho à minha mãe Yudi, ao meu pai Ricardo e ao meu irmão João, que sempre estiveram ao meu lado em todos estes anos, bem como a todos os meus queridos familiares. Deixo um especial agradecimento ao meu avô Roque Pergher (*in memoriam*) pela incansável dedicação na transmissão dos valores do amor ao conhecimento, à cultura, à arte e ao trabalho.

Agradeço a todas as instituições que, direta ou indiretamente, financiaram parte desta trajetória, mesmo não sendo eu bolsista, em especial à CAPES, ao CNPq, à Universidade Federal do Rio Grande do Sul e à FAPESP.

Obrigado a todos os companheiros da Quantitas Asset Management pelos últimos anos de intensa vivência no mercado financeiro, em especial aos times quantitativo e multimercado, com os quais muito aprendi e aprendo todos os dias.

Por mais que um parágrafo não seja suficiente, deixo aqui a minha máxima expressão de amor e gratidão pela companhia da minha maravilhosa Yasmin Noronha, que tem sido, cada dia mais, aquela que me sustenta nos momentos de adversidade, além de minha grande companheira e maior debatedora de ideias nos mais variados temas; que sorte tenho eu em tê-la. Agradeço também aos meus sogros pela presença tão próxima que tivemos nos últimos anos.

ABSTRACT

Portfolio allocation methods based on hierarchical clustering, such as Hierarchical Risk Parity (HRP), avoid the instability of covariance matrix inversion but remain sensitive to the estimation noise that contaminates sample correlations, especially in large, volatile, and sector-concentrated equity universes. This dissertation proposes and validates two allocation methods designed for such conditions, using the Brazilian stock market as the empirical setting. The first method, Orthogonal Hierarchical Risk Parity (OHRP), projects asset returns into a lower-dimensional orthogonal subspace that preserves the local neighborhood structure of the data before the hierarchical clustering and risk allocation stages, with projection parameters optimized in-sample at every rebalancing period. The second method, Sectoral OHRP (S-OHRP), generalizes OHRP by restricting the locality graph to intra-sector connections, thereby injecting an economic prior into the learned representation, and by unifying the in-sample selection into a $\delta \cdot F$ framework that supports the minimization or maximization of arbitrary portfolio objectives, including the Sharpe ratio. This work also introduces the Sectoral Gini Index (SGI), a concentration measure defined over sector-level aggregate weights. The methods are evaluated out of sample against equally weighted, risk parity, and HRP benchmarks, on daily adjusted closing prices of 158 Brazilian stocks spanning January 2011 to February 2026, across estimation windows from 0.2 to 5 years, under a unified backtesting protocol with transaction costs and Wilcoxon signed-rank testing. OHRP achieves statistically significant lower volatility than all benchmarks at every window length, and smaller maximum drawdowns and Pain Index at nearly all, together with the best composed Sharpe ratios at eight of the nine long windows. S-OHRP delivers the best risk-adjusted performance of the entire study, with the highest composed Sharpe ratio (0.482, at a 1.2-year window) and the best average across windows from 0.2 to 2 years, evidence that the sectoral prior regularizes the locality topology without destroying signal. The SGI reveals that the orthogonal methods, most markedly S-OHRP, concentrate allocations within few sectors despite asset-level diversification, a dimension of risk invisible to classical concentration measures.

Keywords: Portfolio allocation; Hierarchical risk parity; Orthogonal projections; Brazilian stock market; Machine learning.

Métodos Ortogonais de Paridade de Risco Hierárquica para Alocação de Carteiras no Mercado de Ações Brasileiro

RESUMO

Métodos de alocação de carteiras baseados em agrupamento hierárquico, como o Hierarchical Risk Parity (HRP), evitam a instabilidade da inversão da matriz de covariância, mas permanecem sensíveis ao ruído de estimação que contamina as correlações amostrais, especialmente em universos acionários grandes, voláteis e setorialmente concentrados. Esta dissertação propõe e valida, no mercado de ações brasileiro, dois métodos de alocação projetados para essas condições. O primeiro método, Orthogonal Hierarchical Risk Parity (OHRP), projeta os retornos dos ativos em um subespaço ortogonal de menor dimensão que preserva a estrutura de vizinhança dos dados antes das etapas de agrupamento hierárquico e alocação de risco, com parâmetros de projeção otimizados in-sample a cada rebalanceamento. O segundo método, Sectoral OHRP (S-OHRP), generaliza o OHRP ao restringir o grafo de localidade a conexões intrassetoriais, injetando assim um prior econômico na representação aprendida, e ao unificar a seleção in-sample em um arcabouço $\delta \cdot F$ que admite a minimização ou maximização de objetivos arbitrários de carteira, incluindo o índice de Sharpe. Este trabalho também introduz o Índice de Gini Setorial (SGI), uma medida de concentração definida sobre os pesos agregados por setor. Os métodos são avaliados fora da amostra contra pesos iguais, paridade de risco e HRP, com preços diários ajustados de 158 ações brasileiras de janeiro de 2011 a fevereiro de 2026, com janelas de estimação de 0,2 a 5 anos, sob um protocolo unificado de backtesting com custos de transação e testes de Wilcoxon. O OHRP alcança volatilidade estatisticamente menor que todos os benchmarks em todas as janelas e drawdowns máximos e Pain Index menores em quase todas, além dos melhores índices de Sharpe compostos em oito das nove janelas longas. O S-OHRP entrega o melhor desempenho ajustado ao risco do estudo, com o maior Sharpe composto (0,482, em 1,2 ano) e a melhor média nas janelas de 0,2 a 2 anos, evidência de que o prior setorial regulariza a topologia de localidade sem destruir o sinal. O SGI revela que os métodos ortogonais, sobretudo o S-OHRP, concentram as alocações em poucos setores apesar da diversificação entre ativos, uma dimensão de risco invisível às medidas clássicas.

Palavras-chave: Alocação de carteiras; Paridade de risco hierárquica; Projeções ortogonais; Mercado de ações brasileiro; Aprendizado de máquina.

LIST OF FIGURES

Figure 5.1 Out-of-sample cumulative returns (net of costs) of OHRP and the benchmark strategies for representative window lengths: $WL = 1.5$ (top), 2.5 (middle), and 4.5 years (bottom). 51
Figure 5.2 Distributions of per-holding-period Sharpe ratio (left) and annualized volatility (right) across the 89 out-of-sample portfolios, per strategy and window length. 52
Figure 5.3 Sectoral composition of the 158-asset universe according to the B3 economic sector classification. 55
Figure 5.4 Out-of-sample cumulative returns (net of costs) of the six strategies for representative window lengths: $WL = 0.40$ (top), 1.20 (middle), and 2.00 years (bottom). 59
Figure 5.5 Distributions of per-holding-period Sharpe ratio (top left), annualized volatility (top right), Gini index (bottom left), and Sectoral Gini Index (bottom right) across the 89 out-of-sample portfolios, per strategy and window length. 60
Figure 5.6 Pairwise correlations between the out-of-sample return series of the six strategies, per window length. 62

LIST OF TABLES

Table 5.1 Composed-series performance measures, OHRP experiment 53
Table 5.2 Consolidated dominance counts, OHRP experiment 53
Table 5.3 Composed-series performance measures, S-OHRP experiment 61
Table 5.4 Consolidated dominance counts, short-window regime 62
Table 5.5 Consolidated dominance counts, long-window regime 62

Table A.1 Individual-portfolio measures, OHRP experiment 76
Table A.2 Individual-portfolio measures, S-OHRP experiment 79

Table B.1 Reference environment of the experiments. 84

LIST OF ABBREVIATIONS AND ACRONYMS

B3	Brasil, Bolsa, Balcão (Brazilian Stock Exchange)
CVaR	Conditional Value at Risk
DD	Maximum Drawdown
EW	Equally Weighted portfolio
GI	Gini Index
HERC	Hierarchical Equal Risk Contribution
HRP	Hierarchical Risk Parity
IBOV	Bovespa index (main stock market index of B3)
LPP	Locality Preserving Projections
OHRP	Orthogonal Hierarchical Risk Parity
OLPP	Orthogonal Locality Preserving Projections
PCA	Principal Component Analysis
PI	Pain Index
RP	Risk Parity
SELIC	Sistema Especial de Liquidação e de Custódia (Brazilian base interest rate)
SGI	Sectoral Gini Index
SMLL	Small Caps (B3 small-capitalization stock index)
S-OHRP	Sectoral Orthogonal Hierarchical Risk Parity
SoR	Sortino Ratio
SR	Sharpe Ratio
SVD	Singular Value Decomposition
TO	Turnover Ratio
WL	Window Length (in-sample estimation window, in years)

LIST OF SYMBOLS

N	Number of available assets
$\boldsymbol{x}_n$	In-sample series of returns of the n-th asset
$\boldsymbol{y}_n$	Projected (lower-dimensional) series of returns of the n-th asset
X	In-sample data matrix ($WL \times N$)
$\boldsymbol{\omega}$	Vector of portfolio weights
Σ	Covariance matrix of asset returns
Σ_{LW}	Ledoit–Wolf shrinkage covariance estimator
ρ	Correlation matrix of asset returns
S	Locality (similarity) matrix
$\tilde{S}$	Sector-constrained locality matrix
D	Diagonal degree matrix of the locality graph
L	Graph Laplacian matrix ($L = D - S$)
W_{PCA}	PCA projection matrix
W_{OLPP}	OLPP orthogonal projection matrix
W	Composed projection matrix ($W = W_{PCA}W_{OLPP}$)
k	Number of nearest neighbors in the locality graph
d	Dimensionality of the projected subspace
r	Fraction of variance retained by the PCA stage
u	Locality bandwidth of the similarity kernel
δ	Optimization direction ($+1$ minimizes, -1 maximizes)
F	Scalar in-sample portfolio objective function
α	Split factor of the recursive bisection
σ_p	Standard deviation of daily portfolio returns
R_f	Daily risk-free rate (SELIC)
Ω_s	Aggregate portfolio weight of sector s
A_s	Set of assets belonging to sector s
c_{bk}	Brokerage (transaction) cost per unit of traded volume

CONTENTS

1 INTRODUCTION 13
1.1 Context and Motivation 13
1.2 Emerging Markets and the Brazilian Stock Market 14
1.3 Problem Statement and Contributions of this Work 15
1.3.1 Goals 15
1.3.2 Main Contributions 16
1.4 Chapters Organization of this Work 18
2 THEORETICAL FOUNDATIONS AND BACKGROUND 19
2.1 Portfolio Allocation Fundamentals 19
2.2 Baseline Allocation Strategies 21
2.2.1 Equally Weighted Portfolio 21
2.2.2 Risk Parity 21
2.3 Hierarchical Risk Parity 22
2.3.1 Hierarchical Tree Clustering 22
2.3.2 Quasi-Diagonalization 24
2.3.3 Recursive Bisection 25
2.4 Dimensionality Reduction and Orthogonal Projections 27
2.4.1 Principal Component Analysis 27
2.4.2 Locality Preserving Projections 27
2.4.3 Orthogonal LPP 28
2.5 Performance and Concentration Metrics 29
2.5.1 Risk-Adjusted Performance 29
2.5.2 Drawdown Measures 30
2.5.3 Concentration and Activity 30
2.6 Statistical Hypothesis Testing 31
3 RELATED WORK 33
3.1 Covariance Estimation and Robust Allocation 33
3.2 Dimensionality Reduction in Finance 34
3.3 HRP Variants and Extensions 34
3.4 Sector Information in Portfolio Construction 35
3.5 Portfolio Allocation in Emerging and Brazilian Markets 36
3.6 Discussion 37
4 PROPOSED ORTHOGONAL HIERARCHICAL RISK PARITY APPROACH 38
4.1 Orthogonal Hierarchical Risk Parity 39
4.1.1 Overview and Rationale 39
4.1.2 Robust Dimensionality Reduction via Orthogonal Locality Preserving Projections 40
4.1.3 The OHRP Algorithm and In-Sample Parameter Optimization 42
4.2 Sectoral Orthogonal Hierarchical Risk Parity 43
4.2.1 Motivation 43
4.2.2 Sector-Constrained Locality Graph 44
4.2.3 Generalized In-Sample Optimization 44
4.2.4 The Sectoral Gini Index (SGI) 45
5 EXPERIMENTAL RESULTS AND DISCUSSION 46
5.1 Orthogonal Hierarchical Risk Parity: Experimental Results 47
5.1.1 Experimental Setup 47
5.1.2 Results and Analysis 48
5.1.2.1 Risk-Adjusted Returns 48

5.1.2.2 Risk Measures 49
5.1.2.3 Concentration and Turnover 50
5.2 Sectoral Orthogonal Hierarchical Risk Parity: Experimental Results 54
5.2.1 Experimental Setup 54
5.2.2 Results and Analysis 54
5.2.2.1 Two Window-Length Regimes 55
5.2.2.2 Risk-Adjusted Returns 56
5.2.2.3 Risk Measures 57
5.2.2.4 Concentration, Turnover, and the SGI 57
5.2.2.5 Cross-Strategy Correlations 58
5.3 Discussion 63
6 CONCLUSIONS AND FUTURE WORK 66
6.1 Publications Based on this Work 70
6.1.1 Publications in Journals 70
6.1.2 Publications in Conferences 70
6.1.3 Manuscripts Under Review 70
REFERENCES 71
APPENDIX A — DETAILED EXPERIMENTAL RESULTS 75
A.1 OHRP Experiment: Individual-Portfolio Measures 76
A.2 S-OHRP Experiment: Individual-Portfolio Measures 79
APPENDIX B — REPRODUCIBILITY OF THE EXPERIMENTS 83
B.1 Code, Data, and Protocol 83
B.2 Reference Environment 83
B.3 Why the Number of Threads Matters 84
APPENDIX C — RESUMO EXPANDIDO EM PORTUGUÊS 85
APPENDIX D — DECLARAÇÃO DE USO DE INTELIGÊNCIA ARTIFICIAL GENERATIVA 89

1 INTRODUCTION

This chapter situates the problem addressed by this dissertation and states what it sets out to accomplish. It begins with the estimation difficulties that motivate hierarchical allocation methods and with the particular features of emerging markets that aggravate them, narrows these into a precise problem statement and a set of goals, summarizes the contributions and the scientific works in which they matured, and closes by outlining how the remainder of the text is organized.

1.1 Context and Motivation

How should capital be divided among a universe of risky assets? The question is among the oldest in quantitative finance, and since the seminal work of Markowitz (1952) the theoretical answer has been, in principle, well understood, namely that portfolio weights should reflect the joint distribution of asset returns, balancing expected reward against risk. In practice, however, that distribution is never observed, since it must be estimated from limited and noisy historical samples, and the optimization machinery tends to amplify whatever errors those estimates contain. Small perturbations in the estimated means and covariances often translate into large swings in the resulting allocation, producing concentrated, unstable portfolios whose out-of-sample performance can be disappointing. Indeed, the naive equally weighted portfolio, which ignores all sample information, has proven remarkably difficult to beat in practice (DeMiguel; Garlappi; Uppal, 2009), which constitutes a sobering benchmark for any optimization-based method.

Hierarchical Risk Parity (HRP) (López de Prado, 2016) offered an appealing way out of this difficulty. Instead of inverting an estimated covariance matrix, HRP organizes the assets into a hierarchy through unsupervised clustering of their correlations, and then distributes risk down that hierarchy by recursive bisection. The approach is robust by design, in that it remains well-defined even when the covariance matrix is singular, and it imposes an economically intuitive tree structure on the allocation. Nevertheless, HRP is not free from estimation problems of its own. Its inputs are still sample correlations, and in large asset universes observed over short windows those correlations are heavily contaminated by noise, so that spurious co-movements can distort the dendrogram on which the entire allocation rests. How to supply the hierarchical machinery with a cleaner and more stable representation of the asset return structure, and how to make that representa-

tion adapt to changing market conditions, is the central question this dissertation seeks to answer.

1.2 Emerging Markets and the Brazilian Stock Market

The question becomes most pressing precisely where portfolio methods are most stressed, which is the case of emerging markets. Relative to their developed counterparts, emerging equity markets tend to exhibit higher and more asymmetric volatility and a greater role of downside risk in expected returns (Bekaert; Harvey, 2003; Estrada, 2002), and their indices are often concentrated in a handful of economic sectors, as is the case of the Brazilian universe studied here (Figure 5.3). These features belong to the data-generating process rather than to the estimator, but they reach the estimator nonetheless, because hierarchical methods take the sample correlation matrix as input, and under heavy tails, asymmetry, and volatility clustering its sampling properties deteriorate, and estimates taken in turbulent periods are biased upward relative to calmer ones (Forbes; Rigobon, 2002). With many assets relative to the length of the estimation window, it is also subject to the high-dimensional noise that the literature on large covariance matrices addresses through regularization (Ledoit; Wolf, 2004a; Bickel; Levina, 2008; Fan; Liao; Mincheva, 2013), as discussed in Section 2.1, and together these conditions place hierarchical methods under considerable strain.

The Brazilian stock market (B3) over the period studied in this dissertation, from January 2011 to February 2026, is a particularly demanding laboratory in this respect. The decade opened in the aftermath of the global commodity boom and soon progressed into the deep domestic recession of 2014–2016, in which a combination of policy missteps and a fiscal crisis produced a fall of about 9% in GDP per capita between 2014 and 2016 (Barbosa Filho, 2017), accompanied by a presidential impeachment and prolonged political instability. What followed was, perhaps surprisingly, a strong equity rally, in which the Bovespa index (IBOV), the main stock market index of the exchange, roughly tripled between early 2016 and early 2020. That rally was abruptly interrupted by the COVID-19 shock of March 2020, when the exchange triggered multiple circuit breakers within days and the index lost more than a third of its value within a few weeks, before staging a rapid but uneven recovery, a market reaction examined in Costa, Silva and Matos (2022). Throughout all of this, the SELIC policy rate oscillated between double-digit levels and historical lows, repeatedly reshaping the opportunity cost of holding equities at

all. An allocation method evaluated over such a period must withstand regime changes, volatility clusters, and sectoral rotations that are far more pronounced than those found in typical developed-market samples. These are exactly the noise conditions under which correlation-based clustering is most fragile, which is why, in our view, the Brazilian market is an ideal testing ground for the methods proposed here. It is worth stressing at the outset that these methods do not attempt to learn a predictive model from such data, which under regime changes of this magnitude would quickly become obsolete. Both are unsupervised, in the sense that they infer a representation of the return structure from the in-sample window alone, without any target variable, and they re-estimate that representation from scratch at every rebalancing (Section 4.1.3).

1.3 Problem Statement and Contributions of this Work

In light of the above, this dissertation addresses the problem of *how to improve the out-of-sample performance of hierarchical risk parity allocation in a large, noisy, sector-concentrated equity universe such as the Brazilian stock market*.

Three specific deficiencies of the existing toolkit are targeted. First, HRP consumes the raw sample correlation structure as-is, with no mechanism to attenuate the noise that distorts its clustering stage. Second, HRP has no tunable parameters, so that, although it re-estimates the correlations at every window, it has no degree of freedom with which to adjust how that structure is represented, and, conversely, any learned representation introduced to help it would bring parameters of its own, requiring a principled, data-driven selection procedure rather than an arbitrary fixed choice. Third, both existing methods and common evaluation practices largely ignore the sectoral dimension of the market, whether as prior information that could inform portfolio construction or as a dimension of concentration risk that asset-level diversification measures simply cannot see.

1.3.1 Goals

The general goal of this dissertation is to develop and empirically validate hierarchical allocation methods based on orthogonal, locality-preserving projections of asset returns, targeted at the Brazilian stock market. This general goal unfolds into the follow-

ing specific goals:

1. To propose OHRP, an extension of HRP in which asset returns are projected into a lower-dimensional orthogonal subspace that preserves the local neighborhood structure of the assets (Eq. (2.11)) before clustering, with projection parameters optimized in-sample at every rebalancing period;
2. To propose S-OHRP, a strict generalization of OHRP that constrains the locality graph to intra-sector connections and generalizes the in-sample selection into a unified $\delta \cdot F$ framework, supporting the minimization or maximization of arbitrary portfolio objectives;
3. To introduce the Sectoral Gini Index (SGI), a concentration measure defined over sector-level aggregate weights, capable of capturing cross-sector concentration that remains invisible to asset-level metrics;
4. To conduct a comprehensive out-of-sample evaluation of the proposed methods against EW, RP, and HRP benchmarks on daily adjusted closing prices of 158 Brazilian stocks spanning January 2011 to February 2026, across estimation windows ranging from 0.2 to 5 years, under a unified backtesting protocol with transaction costs and formal statistical testing.

1.3.2 Main Contributions

The contributions of this dissertation matured incrementally across three scientific works, each resolving a limitation exposed by its predecessor. The research line opened with an exploratory study presented at the IEEE Symposium Series on Computational Intelligence, in Trondheim, Norway (Pergher; Soldera; Scharcanski, 2025), in which the central intuition of this work was first articulated, namely that the dendrogram on which hierarchical allocation rests can be materially improved if the return series are mapped, before clustering, into a lower-dimensional orthogonal subspace that preserves their locality, that is, their local neighborhood structure (Eq. (2.11)). That paper established the viability of the idea, but did so on a restricted experimental setting and with projection parameters largely fixed a priori.

The first contribution of this dissertation, OHRP, grew out of that prototype and is reported in Pergher, Soldera and Scharcanski (2026b). OHRP couples a two-stage projection, PCA followed by orthogonal locality-preserving projections, to the canonical HRP

pipeline and, crucially, treats the projection parameters not as fixed design choices but as quantities re-estimated in-sample at every rebalancing period by exhaustive grid search. This converts a static preprocessing step into an adaptive one, allowing the representation to track changing market conditions instead of presuming a stationary correlation structure. The evaluation reported there showed that the resulting allocations attain systematically lower volatility and smaller maximum drawdowns than HRP, RP, and EW, with statistical support, while at the same time leaving open two structural weaknesses of the method: a locality graph left unconstrained, which admits spurious cross-sector neighbors whenever data are scarce, and an in-sample criterion restricted to variance minimization, which offers no way of expressing a preference for risk-adjusted return.

Those two weaknesses motivate the second contribution, S-OHRP, developed in Pergher, Soldera and Scharcanski (2026a) and currently under peer review. S-OHRP generalizes OHRP along two complementary axes. First, the locality graph is restricted to intra-sector connections, which injects an economic prior into the learned representation and acts as a structural regularizer designed for the conditions in which estimation noise is most damaging, such as short estimation windows. Second, the in-sample selection is unified into a $\delta \cdot F$ framework, in which δ sets the direction of the search and F is an arbitrary scalar portfolio objective, so that criteria such as the Sharpe ratio become admissible alongside variance. Because removing the sectoral restriction recovers OHRP exactly, the generalization is strict, and the comparison between the two methods therefore constitutes a natural ablation. That same work introduces the third contribution, the Sectoral Gini Index, a concentration measure computed over sector-level aggregate weights, which exposes a dimension of concentration risk that asset-level measures are structurally unable to detect.

Considered as a whole, however, this dissertation amounts to more than the sum of those three works. Its distinctive contribution is the unified treatment that renders them comparable, since all experiments were re-executed under a single protocol, sharing the same asset universe, evaluation period, rebalancing frequency, transaction cost model, and statistical methodology, so that OHRP and S-OHRP could be assessed side by side on equal terms rather than through the incommensurable setups of their original publications. The long-window OHRP experiments of Section 5.1 therefore update, rather than reproduce, the originally published figures, although, as will be discussed in due course, every qualitative finding of the published article is preserved under the stricter protocol, which is itself evidence of the robustness of the method. This unified evaluation further

permits a question that neither paper could address in isolation, namely which method best serves which investor objective over the windows on which both were evaluated. The answer, OHRP for risk control, S-OHRP for risk-adjusted return, and S-OHRP$_{\text{Vol}}$ in between, is the practical decision rule of Chapter 6. The full bibliographic records of the three works are collected in Section 6.1.

1.4 Chapters Organization of this Work

The remainder of this dissertation is organized as follows. Chapter 2 reviews the theoretical foundation upon which the contributions are built: the portfolio allocation problem, the benchmark strategies, the HRP algorithm, locality-preserving orthogonal projections, and the performance metrics and statistical tests used throughout the experiments. Chapter 3 surveys the related literature and identifies the gaps this work sets out to fill. Chapter 4 presents the proposed approach, developing first the OHRP method and then its sectoral generalization, S-OHRP, together with the SGI metric. Chapter 5 reports the out-of-sample evaluation of both methods, over long estimation windows (1 to 5 years) for OHRP and short and intermediate windows (0.2 to 2 years) for S-OHRP, and discusses the two sets of results jointly. Finally, Chapter 6 summarizes the findings, discusses the practical implications and limitations of the proposed methods, outlines directions for future research, and records the scientific production derived from this work.

Four appendices close the text. Appendix A reports the individual-portfolio statistics that underlie the analyses of Chapter 5, with the sample mean of each measure and the half-width of its 95% confidence interval. Appendix B lists what is needed to reproduce the experiments exactly, including the numerical setting on which exact reproduction depends. Appendix C presents an extended abstract in Portuguese, as required for dissertations written in English. Appendix D states how generative artificial intelligence tools were used in the preparation of this text.

2 THEORETICAL FOUNDATIONS AND BACKGROUND

Before presenting the proposed methods, it is useful to establish the concepts they build upon. This chapter begins with the portfolio allocation problem itself and the estimation issues that motivate hierarchical approaches (Section 2.1), and then describes the two classical benchmark strategies used throughout the experiments (Section 2.2). Section 2.3 presents the Hierarchical Risk Parity algorithm in some detail, since its three stages are reused verbatim by the methods of Chapter 4. Section 2.4 then reviews dimensionality reduction, with emphasis on the locality-preserving orthogonal projections that constitute the mathematical core of this work. Finally, Sections 2.5 and 2.6 define the out-of-sample performance measures and the statistical methodology adopted in the experimental chapters.

2.1 Portfolio Allocation Fundamentals

Consider, at a given rebalancing date, a universe of N eligible assets, and, since assets enter and leave the investable set, $N = N(t)$ varies over time, and all quantities below refer to the assets eligible at that date. Let $r_i(t)$ denote the daily return of asset i at time t. A portfolio is described by a weight vector $\boldsymbol{\omega} = [\omega_1, \ldots, \omega_N]^T$, with $\sum_i \omega_i = 1$ and, in this work, $\omega_i \geq 0$ (long-only allocations), and its variance is $\sigma_p^2 = \boldsymbol{\omega}^T \Sigma \, \boldsymbol{\omega}$, where Σ is the covariance matrix of asset returns. The target weights $\boldsymbol{\omega}(\tau_m)$ are set at each rebalancing date $\tau_m \in \mathcal{R} = \{\tau_1, \tau_2, \ldots\}$, and the positions are then held without trading until the next rebalancing date τ_{m+1}, so that, for $\tau_m \leq t < \tau_{m+1}$, the weights drift mechanically with relative prices,

$$\omega_i(t) = \frac{\omega_i(\tau_m) \prod_{s=\tau_m+1}^{t} \big(1 + r_i(s)\big)}{\sum_j \omega_j(\tau_m) \prod_{s=\tau_m+1}^{t} \big(1 + r_j(s)\big)}, \qquad r_p(t) = \sum_i \omega_i(t-1)\, r_i(t), \tag{2.1}$$

where $r_p(t)$ is the daily portfolio return.

The modern treatment of this problem originates in the mean-variance framework of Markowitz (1952), which selects $\boldsymbol{\omega}$ by trading off expected return against variance. Despite its elegance, the mean-variance solution requires inverting an estimate of Σ, and its weights are notoriously sensitive to estimation error, to the point that small perturbations in the estimated inputs can produce large swings in the optimal allocation, frequently leading to concentrated and unstable portfolios that underperform out of sample (DeMiguel;

Garlappi; Uppal, 2009). The problem worsens as N grows relative to the length of the estimation window, since the sample covariance matrix then becomes ill-conditioned or even singular. Broadly speaking, three families of remedies have emerged. The first improves the covariance estimate itself. Shrinkage estimators blend the sample covariance with a structured target and guarantee a well-conditioned estimate (Ledoit; Wolf, 2004a); factor models impose a low-dimensional structure on the covariance, explaining co-movements through a small number of common risk factors (Fama; French, 1993); and random matrix theory provides a principled way to isolate the noise in the eigenvalue spectrum, since the eigenvalues of an empirical correlation matrix that fall within the bounds predicted by the Marchenko–Pastur law for purely random data (Marchenko; Pastur, 1967) are statistically indistinguishable from noise and can be filtered out before optimization (Laloux et al., 1999; Bun; Bouchaud; Potters, 2017). The second family regularizes the expected returns, which are the noisiest inputs of all, and its best-known representative is the Black–Litterman model, which combines the market equilibrium with the investor's subjective views in a Bayesian manner (Black; Litterman, 1992). The third avoids the inversion altogether by imposing structure on the allocation rule. This is the route taken by risk parity, which equalizes the risk contributions of the assets and dispenses with expected returns entirely, and by the hierarchical methods that concern us here.

A practical consideration that permeates all experiments in this dissertation is the cost of trading. Whenever the portfolio is rebalanced, the net daily return discounts transaction costs in proportion to the traded volume:

$$r'_p(t) = r_p(t) - c_{bk} \sum_{i=1}^{N} \left| \omega_{p,i}^{bf}(t) - \omega_{p,i}^{af}(t) \right|, \tag{2.2}$$

where ω^{bf} are the drifted weights of Eq. (2.1) at the end of the previous holding period, ω^{af} are the new target weights, and c_{bk} is the brokerage cost, set to 3 basis points throughout this work. It follows that strategies that trade more aggressively must generate enough gross performance to overcome their own turnover, a point that will resurface repeatedly in the experimental chapters.

2.2 Baseline Allocation Strategies

Two allocation rules serve as reference points throughout this dissertation. Both are deliberately simple, and their simplicity is precisely what makes them demanding benchmarks, because neither estimates expected returns and neither is therefore exposed to the estimation error that most penalizes optimized portfolios. They differ in what they equalize, capital in one case and risk contribution in the other, and this difference is enough to place them at opposite ends of the concentration spectrum against which the proposed methods are later compared.

2.2.1 Equally Weighted Portfolio

The equally weighted (EW) portfolio simply assigns uniform weights, $\omega_i^{EW}(t) = 1/N(t)$, where $N(t)$ is the number of tradable assets at time t. One might expect such a rule, which ignores all sample information, to be easily outperformed, yet the opposite is closer to the truth, as DeMiguel, Garlappi and Uppal (2009) show that none of fourteen optimized models consistently beats the $1/N$ rule out of sample, precisely because EW is immune to estimation error. For this reason, EW serves in this work both as a performance benchmark and, being perfectly diversified at the asset level by construction, as the natural reference point for concentration measures.

2.2.2 Risk Parity

Before turning to the rule itself, it is worth recalling where the idea comes from. Modern portfolio theory, inaugurated by Markowitz (1952) and outlined in Section 2.1, established what remains its organizing insight, namely that what deserves to be diversified is risk rather than capital, and because assets are imperfectly correlated a portfolio can be less risky than the weighted sum of its parts. The mean-variance solution operationalizes that insight by trading expected return against variance, which is also where it becomes fragile, since expected returns are the noisiest inputs an allocator has to estimate. Risk parity can be read as an attempt to keep the insight while discarding the requirement, since it retains the covariance structure that governs diversification and abandons the expected-return vector altogether, replacing the goal of maximizing risk-adjusted re-

turn with the more modest one of balancing how much each asset contributes to total risk. The same renunciation, pushed one step further, underlies the hierarchical methods of Section 2.3, which additionally dispense with inverting the covariance matrix.

The risk parity (RP) portfolio (Maillard; Roncalli; Teïletche, 2010) takes a different view and, instead of capital, equalizes the contribution of each asset to overall portfolio risk. Defining the risk contribution of asset i as $RC_i = \omega_i \, \partial\sigma_p / \partial\omega_i$, the RP weights solve

$$\boldsymbol{\omega}_{RP} = \underset{\boldsymbol{\omega}}{\operatorname{argmin}} \; \boldsymbol{\omega}^T \Sigma_{LW} \, \boldsymbol{\omega}, \qquad \text{s.t.} \quad RC_i = \sigma_p / N \;\; \forall i, \tag{2.3}$$

where Σ_{LW} is the linear shrinkage estimator of Ledoit and Wolf (2004b), the asymptotically optimal convex combination of the sample covariance matrix and a scaled identity target, which guarantees a well-conditioned, invertible estimate. By relying on volatility and correlation information while dispensing with mean estimates, which are the noisiest of all inputs, RP typically produces low-concentration portfolios tilted towards low-volatility assets.

2.3 Hierarchical Risk Parity

Hierarchical Risk Parity (HRP), introduced by López de Prado (2016), replaces the quadratic optimization of the Markowitz framework with a three-stage procedure that exploits the hierarchical structure naturally present in asset correlations. Notably, it requires neither the inversion nor even the positive-definiteness of the covariance matrix. Since the methods proposed in this dissertation preserve these three stages intact and intervene only on their inputs, each stage is described below in some detail, accompanied by its pseudocode.[1] Throughout, the N assets are represented by their series of returns, from which the sample covariance matrix $\Sigma \in \mathbb{R}^{N \times N}$ is computed.

2.3.1 Hierarchical Tree Clustering

The first stage organizes the assets into a tree-like structure $\mathcal{T}$ of clusters, grouping together assets whose returns behave similarly over time. From the covariance matrix Σ,

[1] The detailed step-by-step formulation of the three HRP stages presented in this section and the accompanying pseudocode (Algorithms 1–3) were originally elaborated by the author in (Pergher; Soldera; Scharcanski, 2026b), and are reproduced here, adapted to the notation of this dissertation, for the sake of a self-contained presentation.

the Pearson correlation matrix ρ is obtained, and from it a distance is defined between each pair of assets,

$$d_{i,j} = \sqrt{\tfrac{1}{2}\,(1 - \rho_{i,j})}, \tag{2.4}$$

so that highly correlated assets are close to one another and negatively correlated assets are far apart. It is a proper distance in the mathematical sense rather than a mere divergence, because, if $\boldsymbol{z}_i$ denotes the return series of asset i standardized to zero mean and unit variance and scaled by $1/(2\sqrt{T})$, T being the number of observations, then $d_{i,j} = \|\boldsymbol{z}_i - \boldsymbol{z}_j\|$, so that d inherits from the Euclidean norm non-negativity, symmetry, and the triangle inequality (Mantegna, 1999; López de Prado, 2016). Strictly, $d_{i,j} = 0$ whenever $\rho_{i,j} = 1$, so it is a metric on the standardized series and a pseudometric on the assets themselves, a distinction of no practical consequence for empirical return data.

Agglomerative clustering is then performed on this distance matrix. Starting from a set of clusters $\mathcal{C}$ in which each asset forms its own singleton cluster with zero within-cluster sum of squares (the leaves of the tree), the two closest clusters are merged recursively until a single cluster remains (the root). Throughout this dissertation, the Ward linkage criterion (Ward, 1963) is adopted, which at each step merges the pair of clusters that minimizes the increase in total within-cluster variance. Given two clusters C_k and C_l, the Ward distance is

$$\Delta_{\text{Ward}}(C_k, C_l) = \frac{|C_k|\,|C_l|}{|C_k| + |C_l|}\,\left\|\bar{\boldsymbol{z}}_{C_k} - \bar{\boldsymbol{z}}_{C_l}\right\|^2, \qquad \bar{\boldsymbol{z}}_C = \frac{1}{|C|}\sum_{a \in C} \boldsymbol{z}_a, \tag{2.5}$$

where $|C|$ denotes the number of assets in cluster C and $\bar{\boldsymbol{z}}_C$ is the centroid of the cluster, computed on the standardized series $\boldsymbol{z}_a$ defined above, in which the correlation distance of Eq. (2.4) is Euclidean. Δ_{Ward} is therefore the increase in the within-cluster sum of squares caused by the merge. In practice it never requires the vectors $\boldsymbol{z}_a$ explicitly, because, starting from $\Delta_{\text{Ward}}(\{a\}, \{b\}) = d_{a,b}^2/2$ for singletons, it is updated after each merge by the Lance–Williams recurrence, which depends only on the distances. Compared with simpler criteria such as single or average linkage, Ward linkage tends to produce compact, balanced trees, and it is a standard choice in the HRP literature. The complete procedure is given in Algorithm 1.

Algorithm 1 Hierarchical Tree Clustering with Ward Linkage.

Require: Covariance matrix $\Sigma \in \mathbb{R}^{N \times N}$ of the N asset return series.
Ensure: Hierarchical tree structure $\mathcal{T}$.
Compute the correlation matrix: $\rho_{i,j} \leftarrow \Sigma_{i,j} / \sqrt{\Sigma_{i,i}\,\Sigma_{j,j}}$
Compute the distance matrix: $d_{i,j} \leftarrow \sqrt{\frac{1}{2}(1-\rho_{i,j})}$
Initialize clusters $\mathcal{C} \leftarrow \{\{1\}, \ldots, \{N\}\}$, one asset per cluster
Initialize $\mathrm{SS}(C_i) \leftarrow 0 \quad \forall\, C_i \in \mathcal{C}$
while $|\mathcal{C}| > 1$ **do**
    Find $(C_i, C_j) = \mathrm{argmin}_{(C_k, C_l)}\, \Delta_{\text{Ward}}(C_k, C_l),\ \forall\, C_k, C_l \in \mathcal{C}$
    Merge into a new cluster $C_{\text{new}} \leftarrow C_i \cup C_j$
    Update $\mathcal{C} \leftarrow (\mathcal{C} \setminus \{C_i, C_j\}) \cup \{C_{\text{new}}\}$
    $\mathrm{SS}(C_{\text{new}}) \leftarrow \mathrm{SS}(C_i) + \mathrm{SS}(C_j) + \Delta_{\text{Ward}}(C_i, C_j)$ ▷ within-cluster sum of squares
    Record the merge in $\mathcal{T}$
return $\mathcal{T}$

2.3.2 Quasi-Diagonalization

The second stage reorders the rows and columns of the covariance matrix so that similar assets are placed adjacently, concentrating the larger covariance values around the main diagonal while the off-diagonal blocks, which link dissimilar groups of assets, carry smaller values. The reorganized matrix $\Sigma' = P\Sigma P^T$, with P a permutation matrix, is said to be quasi-diagonal. The permutation changes no covariance, and Σ' is no closer to diagonal than Σ, since what it changes is where the large covariances sit, and it is this arrangement that the allocation stage exploits, as explained next.

The reordering is not arbitrary, since it must respect the hierarchy encoded in $\mathcal{T}$, in the sense that assets belonging to the same subtree remain contiguous. Formally, let $\mathcal{T}$ have leaves $\mathbf{z} = \{z_1, \ldots, z_N\}$ and internal nodes $\mathbf{v} = \{v_1, \ldots, v_{N-1}\}$, and let Φ be the space of all admissible orderings of the leaves consistent with the tree (there are 2^{N-1} of them, since at each internal node the left and right branches may be swapped). For a given ordering $\phi \in \Phi$, the sum of the affinities between adjacent leaves is

$$D_{\mathcal{T},\phi} = \sum_{i=1}^{N-1} \Sigma\Big(z_i^{\phi},\ z_{i+1}^{\phi}\Big)\,, \tag{2.6}$$

and the goal is to find the ordering $\phi_{\text{Opt}} = \mathrm{argmax}_{\phi \in \Phi}\, D_{\mathcal{T},\phi}$ that maximizes it. Rather than enumerating all 2^{N-1} orderings, the optimal ordering is found efficiently by the dynamic-programming scheme of Bar-Joseph, Gifford and Jaakkola (2001), which exploits the recursive structure of the tree, in which, for an internal node v with left and right children,

the best achievable adjacency score is

$$M(v) = M(v_{\text{left}}) + M(v_{\text{right}}) + \Sigma_{z^{RM}_{v_{\text{left}}},\, z^{LM}_{v_{\text{right}}}}, \tag{2.7}$$

where $z^{RM}_{v_{\text{left}}}$ is the rightmost leaf of the left branch and $z^{LM}_{v_{\text{right}}}$ the leftmost leaf of the right branch, so that the only interaction between the two branches is the affinity across their shared boundary. Solving $M(\cdot)$ bottom-up from the leaves to the root v_1 yields $M(v_1) = D_{\mathcal{T},\phi_{\text{Opt}}}$ and the optimal ordering itself, as detailed in Algorithm 2.

Algorithm 2 Quasi-Diagonalization via Optimal Leaf Ordering.

Require: Covariance matrix $\Sigma \in \mathbb{R}^{N\times N}$; hierarchical tree $\mathcal{T}$ with internal nodes $v_1, \ldots, v_{N-1}$.
Ensure: Quasi-diagonal covariance matrix Σ'.

function M(v)
  if $|v| = 2$ **then** ▷ leaf-level node with two assets A and B
    return $\Sigma_{A,B}$
  else
    $v_{\text{left}}, v_{\text{right}} \leftarrow$ left and right children of v
    $M^* \leftarrow -\infty$; $\phi^* \leftarrow \{\}$
    for all orderings $\phi^{v_{\text{right}}}$ of the leaves of v_{right} **do**
      for all orderings $\phi^{v_{\text{left}}}$ of the leaves of v_{left} **do**
        $z^{RM} \leftarrow$ rightmost leaf in $\phi^{v_{\text{left}}}$; $z^{LM} \leftarrow$ leftmost leaf in $\phi^{v_{\text{right}}}$
        $M' \leftarrow \mathbf{M}(v_{\text{left}}) + \mathbf{M}(v_{\text{right}}) + \Sigma_{z^{RM},z^{LM}}$
        if $M' > M^*$ **then**
          $M^* \leftarrow M'$; $\phi^* \leftarrow \phi^{v_{\text{left}}} \cup \phi^{v_{\text{right}}}$
    return ϕ^*
$\phi_{\text{Opt}} \leftarrow \mathbf{M}(v_1)$ ▷ evaluate on the root node
$\Sigma' \leftarrow \Sigma_{\phi_{\text{Opt}}}$ ▷ reorder Σ according to ϕ_{Opt}
return Σ'

2.3.3 Recursive Bisection

The third stage assigns the weights by allocating risk down the hierarchy. It rests on the fact that, if the covariance matrix were exactly diagonal, the minimum-variance portfolio would allocate to each asset in inverse proportion to its variance,

$$\omega_i = \frac{1/\sigma_i^2}{\sum_{j=1}^{N} 1/\sigma_j^2}. \tag{2.8}$$

HRP applies this rule to the two halves of a contiguous split of the ordered list, as if they were two uncorrelated assets, which amounts to neglecting the covariances between

the halves, that is, to treating Σ' as approximately block-diagonal along the tree ordering (Cotton, 2024). The quasi-diagonal ordering makes this approximation more plausible, since the blocks neglected at each split are those linking dissimilar groups, and the rule is applied recursively as a divide-and-conquer scheme. All weights are initialized as $\omega_i = 1$. The ordered asset list is split into two contiguous halves, S_{left} and S_{right}, and the variance of each half is computed as that of its inverse-variance-weighted sub-portfolio, that is, for a sub-cluster C with covariance sub-matrix Σ_C,

$$\tilde{\boldsymbol{\omega}}_C = \frac{\operatorname{diag}(\Sigma_C)^{-1}}{\operatorname{tr}\left[\operatorname{diag}(\Sigma_C)^{-1}\right]}, \qquad \tilde{\sigma}_C^2 = \tilde{\boldsymbol{\omega}}_C^T \, \Sigma_C \, \tilde{\boldsymbol{\omega}}_C. \tag{2.9}$$

A split factor is then computed from the variances of the two halves,

$$\alpha = 1 - \frac{\tilde{\sigma}_{\text{left}}^2}{\tilde{\sigma}_{\text{left}}^2 + \tilde{\sigma}_{\text{right}}^2}, \tag{2.10}$$

and the weights of all assets in S_{left} are multiplied by α, while those in S_{right} are multiplied by $1 - \alpha$, so that the less volatile half receives the larger share. For instance, if the left and right halves have annualized variances of 6 and 4, then $\alpha = 0.4$ and $1 - \alpha = 0.6$, and 60% of the risk budget is directed to the less volatile right half. The procedure recurses on each half until every sub-list contains a single asset, at which point $\sum_i \omega_i = 1$ and the final weights are obtained. Algorithm 3 states the procedure.

Algorithm 3 Recursive Bisection for Hierarchical Risk Parity.

Require: Quasi-diagonal covariance matrix Σ'; ordered asset list from $\mathcal{T}$.
Ensure: Portfolio weights $\boldsymbol{\omega}$.

function RECURSIVEBISECTION(cluster, $\boldsymbol{\omega}$, Σ')
  if cluster is a single asset **then**
    return $\boldsymbol{\omega}$ ▷ base case: leaf node
  else
    $S_{\text{left}}, S_{\text{right}} \leftarrow$ left and right sub-clusters
    $\tilde{\sigma}_{\text{left}}^2 \leftarrow \tilde{\boldsymbol{\omega}}_{\text{left}}^T \, \Sigma'_{\text{left}} \, \tilde{\boldsymbol{\omega}}_{\text{left}}$; $\tilde{\sigma}_{\text{right}}^2 \leftarrow \tilde{\boldsymbol{\omega}}_{\text{right}}^T \, \Sigma'_{\text{right}} \, \tilde{\boldsymbol{\omega}}_{\text{right}}$
    $\alpha \leftarrow 1 - \tilde{\sigma}_{\text{left}}^2 / (\tilde{\sigma}_{\text{left}}^2 + \tilde{\sigma}_{\text{right}}^2)$ ▷ variance-ratio split
    $\omega_i \leftarrow \alpha \, \omega_i, \ \forall \, i \in S_{\text{left}}$; $\omega_j \leftarrow (1 - \alpha) \, \omega_j, \ \forall \, j \in S_{\text{right}}$
    RECURSIVEBISECTION($S_{\text{left}}, \boldsymbol{\omega}, \Sigma'_{\text{left}}$)
    RECURSIVEBISECTION($S_{\text{right}}, \boldsymbol{\omega}, \Sigma'_{\text{right}}$)
    return $\boldsymbol{\omega}$
$\boldsymbol{\omega} \leftarrow \{1, \ldots, 1\}$ ▷ initialize weights
$\boldsymbol{\omega} \leftarrow$ RECURSIVEBISECTION($\mathcal{T}_{\text{root}}, \boldsymbol{\omega}, \Sigma'$)

Because the only inputs consumed are the pairwise correlations (for the tree) and the variances (for the bisection), HRP remains well-defined even when the covariance

matrix is singular. Its main fragility, which is also the point of departure for this dissertation, lies elsewhere, in the fact that both the tree and the resulting ordering are estimated from the sample correlation matrix, which in large universes and short windows is heavily contaminated by noise.

2.4 Dimensionality Reduction and Orthogonal Projections

Dimensionality reduction techniques map high-dimensional data into a lower-dimensional subspace that preserves properties of interest. In the context of this dissertation, the objects being projected are the N in-sample series of asset returns, and the property one wishes to preserve is the local similarity structure among them, since this is precisely the information the hierarchical clustering stage consumes.

2.4.1 Principal Component Analysis

Principal Component Analysis (PCA) finds the orthogonal directions of maximal variance of the data by solving the eigenproblem $XX^T\vec{v} = \lambda\vec{v}$, where X is the (centered) data matrix. Projecting onto the leading eigenvectors yields the best low-rank reconstruction of the data in the mean-squared sense, and discarding the components with the smallest eigenvalues acts, in addition, as a denoising device, since in financial data the trailing principal components tend to be dominated by idiosyncratic noise. PCA is, however, a fundamentally global method, which preserves overall variance but has no mechanism to retain the neighborhood relations among individual data points. In this work the data points are the assets, each represented by its return series over the estimation window, and two assets are neighbors when these series are close in Euclidean distance, a proximity that reflects both how strongly the assets co-move and how similar their volatilities are (Section 4.1.2), so that neighborhood is cross-sectional, not temporal.

2.4.2 Locality Preserving Projections

Locality Preserving Projections (LPP) (He; Niyogi, 2004) address precisely this limitation. LPP builds an adjacency graph over the data points, here the N assets, each a point in $\mathbb{R}^{WL}$ given by its return series $\boldsymbol{x}_n$ over the estimation window (the columns

of the $WL \times N$ matrix X). This is the transpose of the more usual view in finance, in which each date is a point in $\mathbb{R}^N$ formed by the cross-section of returns, and it is what makes the graph a graph of assets. LPP then weights its edges by a similarity kernel, and seeks the linear projection that preserves the *locality* of the data, that is, its neighborhood structure as encoded in the graph, so that points that are neighbors in the original space remain close after projection. Formally, let $S \in \mathbb{R}^{N\times N}$ be the similarity matrix of the graph, whose entry $S_{nn'}$ is the edge weight between assets n and n' (the indices run over assets, not over observations), let $D \in \mathbb{R}^{N\times N}$ be the diagonal degree matrix, whose entry $D_{nn} = \sum_{n'} S_{nn'}$ is the total similarity of asset n to its neighbors, and let $L = D - S \in \mathbb{R}^{N\times N}$ be the graph Laplacian, so that XLX^T and XDX^T are $WL \times WL$ matrices. LPP then solves

$$\begin{aligned} &\min_{W} \sum_{n,n'} \left\| W^T \boldsymbol{x}_n - W^T \boldsymbol{x}_{n'} \right\|^2 S_{nn'} = \min_{W} \operatorname{tr}\left(W^T X L X^T W\right) \\ &\text{subject to} \quad W^T X D X^T W = I, \end{aligned} \tag{2.11}$$

so that pairs of neighbors (large $S_{nn'}$) are penalized for being mapped far apart, while non-neighbors do not enter the objective. The solution is given by the eigenvectors of the generalized eigenproblem involving XLX^T and XDX^T associated with the smallest eigenvalues. PCA and LPP preserve different properties, global variance and local neighborhoods, and their sensitivity to outlying assets differs accordingly. In PCA an asset far from the others enters the variance with quadratic weight, whereas in LPP built on a k-nearest-neighbor graph with a heat kernel it enters only through its few edges, whose weights vanish with distance. This protection is a property of the graph construction and of the Euclidean heat kernel rather than of locality as such, and it need not hold for other choices of distance or weighting.

2.4.3 Orthogonal LPP

The basis produced by LPP, however, is not orthogonal in general, which distorts metric relations in the projected space and complicates interpretation. Orthogonal Locality Preserving Projections (OLPP) (Cai et al., 2006) enforce orthogonality by constructing the basis iteratively, so that each new direction solves the LPP-type eigenproblem restricted to the orthogonal complement of the directions already found. The result is an orthonormal basis that still preserves locality, combining the neighborhood-

preserving graph construction of LPP with the geometric benefits of orthogonal representations. OLPP has been successfully applied in other domains where preserving the local structure of high-dimensional data is essential, among them face recognition, for which Soldera, Behaine and Scharcanski (2015) employed a customized OLPP with soft-margin maximization, and it was precisely this prior application that inspired the use of OLPP as the dimensionality-reduction core of the methods proposed in this dissertation. In a portfolio context, the orthogonal representation means that overlapping sources of risk are isolated into independent dimensions, mitigating multicollinearity among the projected series. The complete OLPP construction, including its composition with a preliminary PCA stage, is presented in Section 4.1.2 as an integral part of the OHRP algorithm.

2.5 Performance and Concentration Metrics

No single number summarizes portfolio quality, and for this reason the out-of-sample evaluation relies on a set of complementary metrics, all computed over the net-of-cost daily returns of Eq. (2.2) with an annualization factor of 252 business days. Throughout, $\bar{R}_p$ denotes the mean daily net portfolio return and R_f the daily risk-free rate, taken in all experiments as the effective SELIC rate, the overnight rate published by the Central Bank of Brazil for each business day, whose rates are compounded into a cumulative index so that $R_f(t)$ is the rate accrued over business day t, on the same close-to-close calendar as the asset returns.

2.5.1 Risk-Adjusted Performance

The annualized Sharpe ratio (Sharpe, 1966) measures the excess return earned per unit of total risk,

$$SR = \sqrt{252}\,\frac{\bar{R}_p - R_f}{\sigma_p}, \tag{2.12}$$

where σ_p is the standard deviation of daily returns. Since investors are arguably indifferent to upside volatility, the Sortino ratio (Sortino; Price, 1994) penalizes only its harmful component, replacing σ_p by its downside counterpart,

$$SoR = \sqrt{252}\,\frac{\bar{R}_p - R_f}{\sigma_p^{down}}, \tag{2.13}$$

with σ_p^{down} computed over negative excess returns only. The annualized volatility $\sigma = \sqrt{252}\,\sigma_p$ and the annualized mean return $\mu = 252\,\bar{R}_p$ are also reported as standalone risk and return proxies.

2.5.2 Drawdown Measures

The drawdown at time t is the relative loss with respect to the running maximum of the cumulative wealth curve $P(t)$,

$$D(t) = \frac{P_{\max}(t) - P(t)}{P_{\max}(t)}, \tag{2.14}$$

and two summaries of it are used in this work: the maximum drawdown $DD = \max_t D(t)$, which captures the worst peak-to-trough loss, and the Pain Index,

$$PI = \frac{1}{T}\int_0^T D(t)\,dt, \tag{2.15}$$

which averages the drawdown over the whole period and therefore captures the *persistence* of losses rather than only their worst realization. Drawdowns deserve this double treatment because investors experience them directly, as unrealized losses, which makes such metrics particularly relevant in volatile markets like the one studied here.

2.5.3 Concentration and Activity

The Gini index, applied to the portfolio weights ($\sum_i \omega_i = 1$), measures how unevenly the capital is allocated,

$$GI = \frac{\sum_{i,j}|\omega_i - \omega_j|}{2N}, \tag{2.16}$$

with $GI = 0$ for the EW portfolio and $GI \to (N-1)/N$ under full concentration in a single asset. Its sectoral extension, the SGI, is one of the contributions of this dissertation, and its presentation is therefore deferred to Section 4.2.4. Finally, trading activity is measured by the turnover ratio at each rebalancing date,

$$TO(t) = \frac{1}{2}\sum_{i=1}^{N}\left|\omega_{p,i}^{bf}(t) - \omega_{p,i}^{af}(t)\right|, \tag{2.17}$$

which quantifies the fraction of the portfolio effectively traded, and connects directly to the cost model of Eq. (2.2).

2.6 Statistical Hypothesis Testing

Comparing strategies solely through composed, full-period statistics can be misleading, because a single fortunate sub-period may dominate the aggregate result, giving an inflated impression of superiority. For this reason, the experimental chapters complement the composed metrics with paired hypothesis tests over the per-holding-period observations. For each pair of strategies evaluated under the same window length, the two samples of 89 paired observations are compared with the Wilcoxon signed-rank test (Wilcoxon, 1945), which ranks the absolute differences between paired observations and verifies whether their distribution is symmetric around zero. Crucially, the test makes no normality assumption, an important property given the heavy-tailed nature of financial returns, although, as discussed below, it is not free of assumptions. A strategy is declared superior to another for a given metric only when the null hypothesis is rejected at the 95% confidence level in the favorable direction, and aggregating these pairwise outcomes across benchmarks and window lengths yields the dominance counts reported in Chapter 5.

The natural parametric alternative is Student's t-test for paired samples, which compares the mean of the paired differences against zero. It was employed alongside the Wilcoxon test in Pergher, Soldera and Scharcanski (2026b), and for completeness both tests are reported in Chapter 5. The t-test, however, assumes that the paired differences are normally distributed, an assumption that daily portfolio returns violate in a well-documented manner, since they are leptokurtic, negatively skewed, and prone to volatility clustering, all the more so in emerging markets. When normality fails, the t-test tends to over-reject the null hypothesis, declaring differences significant that a distribution-free procedure would not endorse. For this reason the Wilcoxon test is adopted as the primary criterion throughout this dissertation, with the t-test reported as a secondary reference, and, as will be seen in Section 5.3, the two procedures agree in most comparisons, and in the OHRP experiment, where they disagree, the pattern is the one this reasoning anticipates.

Neither test is assumption-free, since both treat the 89 paired differences as independent, and under the null the Wilcoxon test also requires them to be symmetric around

zero and the t-test, normal. Since holding periods do not overlap, consecutive differences are nearly uncorrelated for return-based measures (lag-one autocorrelation of about 0.06 for the Sharpe ratio), but volatility clustering carries over to the risk measures (about 0.23 for volatility, an effective sample of about 55). Their p-values are thus approximate, and a dependence-robust alternative would be the bootstrap test of Ledoit and Wolf (2008).

3 RELATED WORK

Having established the theoretical background, we now turn to the literature that surrounds this dissertation, with two purposes in mind: to acknowledge the ideas this work builds upon, and to delimit precisely where it departs from them. The discussion, which positions this work rather than surveying each field exhaustively, is organized around the three ingredients of the proposed methods, namely robust allocation under estimation error (Section 3.1), dimensionality reduction in finance (Section 3.2), and hierarchical allocation schemes (Section 3.3), followed by the two dimensions that are specific to this work: the use of sector information (Section 3.4) and the empirical evidence for emerging and Brazilian markets (Section 3.5). Section 3.6 closes the chapter by summarizing the gaps this dissertation sets out to fill.

3.1 Covariance Estimation and Robust Allocation

The sensitivity of optimized portfolios to estimation error is among the most extensively documented phenomena in the portfolio-choice literature. In a particularly influential study, DeMiguel, Garlappi and Uppal (2009) show that estimation windows of unrealistic length would be required for sample-based mean-variance portfolios to reliably outperform the naive $1/N$ rule, a result that established EW as the de facto robustness benchmark against which any optimization-based method must be measured.

A first family of remedies attacks the covariance estimate itself. Shrinkage estimators (Ledoit; Wolf, 2004a) combine the sample covariance matrix with a structured target, guaranteeing well-conditioned estimates even in large universes, an idea later extended to nonlinear shrinkage of each eigenvalue (Ledoit; Wolf, 2017; Ledoit; Wolf, 2020), whereas, in high dimensions, banding, thresholding, and factor-based estimators exploit sparsity or low-rank structure instead (Bickel; Levina, 2008; Fan; Liao; Mincheva, 2013). A complementary line of research draws on random matrix theory, starting from Laloux et al. (1999), who demonstrated that the bulk of the eigenvalue spectrum of empirical financial correlation matrices is statistically indistinguishable from pure noise, a finding that motivated a whole family of denoising procedures that filter the spectrum before optimization, reviewed in Bun, Bouchaud and Potters (2017). These approaches share their diagnosis with the present work, namely that the sample correlation structure is heavily noise-contaminated, but they differ in where they intervene, since they act

on the input matrix while keeping the quadratic optimization intact. The methods proposed in this dissertation act instead on the *representation* of the data, through orthogonal locality-preserving projections, and on the allocation rule itself, through hierarchical risk allocation, thereby avoiding covariance inversion altogether.

3.2 Dimensionality Reduction in Finance

Projecting asset returns onto a reduced set of components is a long-standing practice in finance, running from the statistical factor models built on principal components (Connor; Korajczyk, 1986) to the latent-factor and machine-learning representations of recent years (Kelly; Pruitt; Su, 2019; Lettau; Pelger, 2020; Gu; Kelly; Xiu, 2020). Among the latter, Gu, Kelly and Xiu (2021) introduced autoencoder-based asset pricing models in which latent factors and their loadings are learned jointly, improving out-of-sample pricing performance relative to linear factor models. Within portfolio construction specifically, dimensionality reduction has typically served one of two purposes: stabilizing the covariance estimate (Fan; Liao; Mincheva, 2013), or extracting return-predictive signals (Avellaneda; Lee, 2010; Gu; Kelly; Xiu, 2020).

The use made of dimensionality reduction in this dissertation differs from both in two respects. First, the projection of choice is locality-preserving and orthogonal (OLPP), selected not for pricing or prediction, but for preserving the similarity structure among assets, which is precisely the information consumed by the hierarchical clustering stage of HRP. Second, and perhaps more importantly, the projection is not fixed once and for all, since its hyperparameters are re-optimized at every rebalancing window against an in-sample portfolio objective, making the representation itself an adaptive, integral component of the allocation strategy rather than a mere preprocessing step.

3.3 HRP Variants and Extensions

Since the introduction of HRP by López de Prado (2016), a growing literature has explored variations of its three-stage design. Raffinot (2017) proposed hierarchical clustering-based asset allocation with alternative linkage criteria and tree-cutting heuristics, an idea later extended into the Hierarchical Equal Risk Contribution (HERC) approach (Raffinot, 2018), in which equal risk contributions are allocated across the clusters

identified by the dendrogram. Subsequent studies have examined the sensitivity of hierarchical allocation to the linkage choice and to the ordering of the assets (Raffinot, 2017; Schwendner et al., 2021), to the distance metric (Salas-Molina et al., 2025), and to the covariance estimator (Jain; Jain, 2019), while others have replaced the clustering stage by graph-based structures (Ciciretti; Pallotta, 2024) or by a fast correlation ranking (Salas-Molina; Nin, 2026), extended the allocation to alternative risk measures (Raffinot, 2018), or recovered the cross-cluster covariance that the bisection discards (Cotton, 2024). The evidence does not single out one stage as the weak link, given that alternative linkages often improve on the single linkage of the original algorithm (Schwendner et al., 2021), yet in synthetic experiments without block structure a random ordering performs almost as well as the hierarchical one (Salas-Molina; Nin, 2026), which places much of the burden on the quality of the correlation estimates fed to the algorithm, precisely where OHRP intervenes.

The common thread running through this family of methods is that they modify the clustering or allocation stages themselves. The approach taken in this dissertation is different in kind, because OHRP (Section 4.1) keeps the three canonical HRP stages intact and intervenes *before* them, supplying the clustering algorithm with the kind of denoised, orthogonal, locality-preserving representation described in Section 3.2, whose parameters are optimized in-sample at every rebalancing. To the best of our knowledge, no prior HRP variant performs dynamic, per-window optimization of a learned representation, let alone constrains that representation with sector information, as S-OHRP does (Section 4.2).

3.4 Sector Information in Portfolio Construction

That firms within the same industry sector share exposure to common risk factors, and consequently exhibit systematically higher intra-sector return correlations, is well established (Moskowitz; Grinblatt, 1999; Chan; Karceski; Lakonishok, 1999; Bhojraj; Lee; Oler, 2003), and industry classification has long been a standard covariate in asset pricing and risk modeling. In portfolio construction, sector information is traditionally employed in one of three ways: through explicit constraints on sector weights, through industry factor models embedded in the covariance estimate (Chan; Karceski; Lakonishok, 1999), or through top-down allocation schemes that first distribute capital across sectors and only then within them. What these approaches have in common is that the sector enters as a constraint or factor acting on the *weights*, and recent work derives industry groupings

from the data rather than from exchange classifications (Hoberg; Phillips, 2016).

The S-OHRP method proposed in this dissertation uses sector information in a structurally different manner, since the sector restricts the topology of the locality graph that defines the subspace projection, acting as a regularizer on the learned representation rather than on the final allocation, in the spirit of the observation that well-chosen constraints act as implicit shrinkage (Jagannathan; Ma, 2003). The portfolio weights themselves remain entirely free of explicit sector constraints, so whether this indirect mechanism also diversifies across sectors is an empirical question, examined with the SGI in Section 5.2.

3.5 Portfolio Allocation in Emerging and Brazilian Markets

Emerging markets exhibit structural features that challenge allocation methods calibrated on developed-market data, such as higher and asymmetric volatility and a greater weight of downside risk in expected returns (Bekaert; Harvey, 2003; Estrada, 2002), besides index concentration in a few sectors, a feature of the Brazilian universe documented in Figure 5.3. For the Brazilian market in particular, the empirical evidence has consistently favored risk-based allocation. Rubesam and Beltrame (2013) showed that minimum-variance portfolios, with covariance matrices estimated by methods ranging from the sample covariance to multivariate GARCH models, delivered higher returns with lower risk than the IBOV, the equally weighted portfolio, and the maximum-Sharpe portfolio. Along similar lines, Caldeira et al. (2017) reported gains from using realized covariance estimators based on intraday data for minimum-variance portfolio selection in Brazil, while Souza et al. (2017) provided an early assessment of risk parity strategies in the Brazilian market.

Closer to the subject of this dissertation, Rubesam (2022) combined machine-learning return forecasts with equal-risk-contribution aggregation in Brazilian equities and found that the forecasts add value, but that long-short strategies suffer from the high volatility of their short legs and that aggregating them by equal risk contributions significantly improves their risk-adjusted performance. Return-driven signals, in other words, pay off in this market when they are channelled through a risk-based allocation, which is also how S-OHRP works, since it selects its projection by in-sample Sharpe ratio and allocates through the risk-based stages of HRP. The S-OHRP results add that such return-driven selection is strong on average but uneven from one window to the next

(Section 5.2). Most directly related to the present work, Reis et al. (2023) conducted an out-of-sample evaluation of HRP in the Brazilian market and found that it outperforms the market index but is generally not the best of the traditional techniques, matching them only in some performance measures. This last result is worth emphasizing, since it delimits the state of the art that this dissertation advances, in which standard HRP, applied as-is to Brazilian data, does not exhibit the systematic risk dominance that OHRP achieves in Section 5.1 under a comparable protocol.

3.6 Discussion

Taken together, the works discussed above leave three gaps, and it is at their intersection that this dissertation operates. First, HRP variants modify the clustering or allocation stages but do not learn a representation of the returns that adapts to the data, that is, one whose parameters are re-selected at every rebalancing from the in-sample window, while denoising approaches filter the covariance matrix but retain inversion-based optimization, and OHRP fills this gap by coupling an orthogonal locality-preserving projection, whose parameters are optimized in-sample at every window, to the intact HRP pipeline. Second, sector information has been used as a weight constraint or covariance factor, but not as a structural prior on the neighborhood graph of a learned projection, and S-OHRP introduces precisely this mechanism, together with a generalized $\delta \cdot F$ objective that admits Sharpe maximization. Third, the Brazilian evidence on hierarchical methods is recent and shows HRP matching traditional techniques in some measures rather than outperforming them, and no existing study measures the cross-sector concentration of allocation strategies, and this dissertation contributes a comprehensive two-regime evaluation on B3 data, and proposes the SGI metric to fill that measurement gap.

4 PROPOSED ORTHOGONAL HIERARCHICAL RISK PARITY APPROACH

This chapter presents the methodological contributions of this dissertation, a family of hierarchical allocation methods built around orthogonal, locality-preserving projections of asset returns. The guiding idea is to change what the HRP machinery sees, rather than modifying the machinery itself. Asset returns are projected into a lower-dimensional orthogonal subspace that preserves the local neighborhood structure of the assets (Eq. (2.11)) before the hierarchical clustering and risk allocation stages take place, and the parameters of that projection are re-optimized in-sample at every rebalancing period, so that the representation adapts to the dependence structure of each estimation window.

The approach is developed in two stages. Section 4.1 introduces Orthogonal Hierarchical Risk Parity (OHRP), originally published in (Pergher; Soldera; Scharcanski, 2026b), in which the locality graph is built from the return series alone and the projection parameters are selected by in-sample volatility minimization. An embryonic version of this idea, in which dynamic orthogonal lower-dimensional projections were first applied to hierarchical risk allocation, was presented as a conference paper at the IEEE Symposium Series on Computational Intelligence (Pergher; Soldera; Scharcanski, 2025), and later matured into the complete method described here. Section 4.2 then presents Sectoral Orthogonal Hierarchical Risk Parity (S-OHRP) (Pergher; Soldera; Scharcanski, 2026a), which generalizes OHRP along two complementary directions: by constraining the locality graph to intra-sector connections, thereby injecting an economic prior into an otherwise purely statistical construction, and by replacing the fixed variance objective with a unified $\delta \cdot F$ selection framework that accommodates arbitrary scalar portfolio objectives. That section also introduces the Sectoral Gini Index (SGI), a novel concentration measure defined over sector-level aggregate weights.

Because setting the sector partition to the full asset universe recovers OHRP exactly, the second method is a strict generalization of the first rather than a competing design. This nesting is what makes the empirical comparison of Chapter 5 informative, because contrasting the two isolates the contribution of each ingredient separately.

4.1 Orthogonal Hierarchical Risk Parity

This section develops the first of the two proposed methods. The presentation proceeds from the motivation to the machinery: why a learned representation should help the clustering stage at all, how the orthogonal locality-preserving projection is constructed, and how its parameters are chosen at each rebalancing without recourse to out-of-sample information.

4.1.1 Overview and Rationale

The HRP algorithm builds a hierarchy of assets from the correlation matrix of their returns and allocates risk through recursive bisection, avoiding the matrix inversion required by classical mean-variance optimization. However, when applied to large universes of multi-sector assets subject to short-term changes in market conditions, the sample correlation matrix becomes noisy, and the resulting clustering hierarchy grows sensitive to spurious co-movements, which, in turn, tends to degrade the out-of-sample performance of HRP-based portfolios.

OHRP addresses this limitation by learning, at every rebalancing date and from the preceding estimation window, an orthogonal lower-dimensional representation of the in-sample asset returns. The projection is designed to preserve the local similarity structure among assets, which is the very information the clustering stage relies on, while attenuating correlations across distinct groups of assets. Intuitively, the orthogonal representation isolates overlapping sources of risk into independent dimensions, mitigating multicollinearity and providing the clustering algorithm with more stable inputs. Besides, OHRP does not commit to a single, fixed projection, since its parameters are selected by a grid search that minimizes the in-sample portfolio volatility, so that the projector itself is re-selected at every rebalancing to fit the dependence structure of the most recent estimation window, that is, its correlations and volatilities and, through them, the neighborhood graph among the assets.

4.1.2 Robust Dimensionality Reduction via Orthogonal Locality Preserving Projections

Let N be the number of available assets and let $\boldsymbol{x}_n$ denote the in-sample series of returns of the n-th asset, so that the in-sample data form a $WL \times N$ matrix X, where WL is the rolling window length in days. The construction starts with a locality graph over the assets, in which the returns series $\boldsymbol{x}_n$ and $\boldsymbol{x}_{n'}$ are assumed to be neighbors if, and only if, $\boldsymbol{x}_{n'}$ is among the k nearest neighbors of $\boldsymbol{x}_n$ or $\boldsymbol{x}_n$ is among the k nearest neighbors of $\boldsymbol{x}_{n'}$, with $k \leq N$. Nearest neighbors are determined by the Euclidean distance $\|\boldsymbol{x}_n - \boldsymbol{x}_{n'}\|$ between the raw daily return series, which are deliberately not rescaled to unit variance, and the implications of this choice are discussed after Eq. (4.1). The similarity matrix S is then defined as

$$S_{nn'} = \begin{cases} e^{-\frac{\|\boldsymbol{x}_n - \boldsymbol{x}_{n'}\|^2}{u}} & \text{if } \boldsymbol{x}_n \text{ and } \boldsymbol{x}_{n'} \text{ are neighbors,} \\ 0 & \text{otherwise,} \end{cases} \tag{4.1}$$

where each element $S_{nn'}$ measures the similarity between assets n and n'. The neighborhood relation is thus symmetrized, so that $S = S^T$. Symmetry is what makes the construction below well posed, because the degree matrix D and the Laplacian $L = D - S$ are symmetric, L is positive semidefinite, and the matrices XDX^T and XLX^T are symmetric, so that the generalized eigenproblem that defines the OLPP basis has real eigenvalues. The locality bandwidth u, which determines how strongly neighboring return series affect each other, is estimated as $u = \bar{Q}$, the mean of the full pairwise Euclidean distance matrix Q computed over all N series of returns.

The choice of the Euclidean distance in Eq. (4.1) deserves a comment, since it is not the only defensible measure of proximity between return series. Because daily mean returns are negligible relative to their dispersion, $\|\boldsymbol{x}_n - \boldsymbol{x}_{n'}\|^2 \approx WL\,(\sigma_n^2 + \sigma_{n'}^2 - 2\rho_{n,n'}\sigma_n\sigma_{n'})$, so that the Euclidean distance between raw returns blends two kinds of information: how strongly two assets co-move, through the correlation $\rho_{n,n'}$, and how similar their volatilities σ_n and $\sigma_{n'}$ are. Had the series been standardized, the expression would reduce to $2\,WL\,(1 - \rho_{n,n'})$, a monotone transformation of the correlation distance of Eq. (2.4) used by HRP (Mantegna, 1999). Other alternatives include the Mahalanobis distance, which accounts for the dependence among observations, and elastic measures such as dynamic time warping, which tolerate lags between series. The Euclidean dis-

tance was retained here for two reasons. It is the native geometry of the heat-kernel locality graph on which OLPP is built (Cai et al., 2006), and it supplies the projection with volatility information that the correlation-based clustering stage of HRP discards, so that the two stages see complementary aspects of the data. A systematic comparison of distance measures for the locality graph is left for future work (Chapter 6).

The final subspace transformation $\boldsymbol{x}_n \to \boldsymbol{y}_n$ is given by

$$\boldsymbol{y}_n = W^T \boldsymbol{x}_n, \qquad W = W_{PCA}\, W_{OLPP}, \tag{4.2}$$

where $W_{PCA} = [\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_{d_{\max}}]$ first projects each series of returns into the PCA orthogonal subspace spanned by the $d_{\max}$ leading eigenvectors, with $d_{\max}$ defined in Eq. (4.3) below. The principal components with smaller eigenvalues are systematically discarded so as to retain a fraction $r \in [0, 1]$ of the variability of the original N components. In practice, the eigenproblem $XX^T\vec{v} = \lambda\vec{v}$ is solved efficiently via the Singular Value Decomposition (SVD) of X, whose singular values are the square roots of the eigenvalues, $\lambda_f = S_{SVD}(f,f)^2$. Since each eigenvalue measures the variance carried by the corresponding component, the maximum admissible dimensionality $d_{\max}$ is the largest number of eigenvectors satisfying

$$\sum_{f=1}^{d_{\max}} S_{SVD}(f,f)^2 \le r \sum_{f=1}^{\min(N,WL)} S_{SVD}(f,f)^2, \tag{4.3}$$

where S_{SVD} is the diagonal matrix of singular values, with $d_{\max} \ge 1$, so that the desired dimensionality satisfies $d \in [1, d_{\max}]$.

The orthogonal locality-preserving basis $[\vec{g}_1, \vec{g}_2, \ldots, \vec{g}_d]$ is obtained as follows. One first computes the diagonal degree matrix D, with $D_{nn} = \sum_{n'} S_{nn'}$, and the Laplacian matrix $L = D - S$. The first basis vector $\vec{g}_1$ is the eigenvector associated with the smallest eigenvalue of $(XDX^T)^{-1}XLX^T$. Defining

$$G^{(p-1)} = [\vec{g}_1, \vec{g}_2, \ldots, \vec{g}_{p-1}], \tag{4.4}$$

$$H^{(p-1)} = [G^{(p-1)}]^T (XDX^T)^{-1} G^{(p-1)}, \tag{4.5}$$

each subsequent vector $\vec{g}_p$, for $2 \le p \le d$, is the eigenvector associated with the smallest

eigenvalue of

$$M^{(p)} = \left\{ I - (XDX^T)^{-1} G^{(p-1)} [H^{(p-1)}]^{-1} [G^{(p-1)}]^T \right\} (XDX^T)^{-1} XLX^T, \tag{4.6}$$

so that orthogonality is enforced by construction, one direction at a time. The resulting OLPP subspace is spanned by $W_{OLPP} = [\vec{g}_1, \vec{g}_2, \ldots, \vec{g}_d]$, and the composed projection $W = W_{PCA} W_{OLPP}$ maps each of the N series of asset returns into its d-dimensional representation $\boldsymbol{y}_n$. These compact representations are what feed the subsequent hierarchical clustering and risk allocation stages.

4.1.3 The OHRP Algorithm and In-Sample Parameter Optimization

At each rebalancing date t, OHRP takes as in-sample data X the returns of the eligible assets over the estimation window formed by the WL days preceding t, and proceeds in three steps. First, these data are projected into the orthogonal subspace defined by the parameters (k, d, r), following the scheme of Section 4.1.2. Second, the projected series $\boldsymbol{y}_n$ feed the three canonical HRP stages, namely hierarchical tree clustering with Ward linkage, quasi-diagonalization of the covariance matrix, and recursive bisection (see Chapter 2), yielding the portfolio weights $\boldsymbol{\omega}$. Third, the candidate portfolio is evaluated on the in-sample data in terms of its realized volatility.

The projection parameters are selected by an exhaustive grid search that minimizes the in-sample variance:

$$(k, d, r)_{\text{opt}} = \underset{(k,d,r)}{\operatorname{argmin}} \; \sigma^2_{\text{in}} \left(\boldsymbol{\omega}(k, d, r) \right), \tag{4.7}$$

with $k \in \{1, 3, 5, 7, 9, 11, 13\}$, $d \in \{5, 10, 15, 20, 25, 30, 35\}$, and $r \in \{0.92, 0.94, 0.96, 0.98, 1.00\}$. In the originally published version of OHRP (Pergher; Soldera; Scharcanski, 2026b), the PCA variance retention was kept fixed at $r = 0.90$, but that fixation was a convenience choice rather than a methodological requirement, and in the unified protocol adopted here r is optimized jointly with k and d. The optimal triple $(k, d, r)_{\text{opt}}$ is then used to assemble the portfolio held during the subsequent out-of-sample period, and the whole procedure repeats with newer in-sample data at every rebalancing date. As will be seen in Section 4.2.3, this in-sample selection scheme admits a natural generalization to arbitrary objective functions.

It should be noted that this in-sample selection is not a learning procedure in the supervised sense. No target variable is predicted, and no fitted quantity is carried from one window to the next, because the projection is unsupervised and the grid search merely selects, among a finite set of candidate projections, the one whose portfolio best meets the in-sample objective. The method thus adapts to changing market conditions by re-estimation at every window, rather than by accumulating a model whose validity would depend on the persistence of past regimes.

4.2 Sectoral Orthogonal Hierarchical Risk Parity

The method just described leaves two questions open: how to protect the locality graph from spurious cross-sector similarities in short estimation windows, and how to let investors express objectives other than volatility minimization. The Sectoral Orthogonal Hierarchical Risk Parity (S-OHRP) method (Pergher; Soldera; Scharcanski, 2026a) answers both at once, and this section develops it.

4.2.1 Motivation

Financial markets exhibit a well-documented sectoral structure, in which firms within the same industry sector share exposure to common risk factors, such as an energy crisis, a shift in global demand for a commodity, or a regulatory change, and, as a result, tend to display higher intra-sector return correlations relative to cross-sector pairs. The OHRP locality graph of Eq. (4.1), however, is entirely blind to this structure, since nearest neighbors are determined purely by the Euclidean distance between return series. In long windows this was expected to matter less, since more data make the estimated similarities more reliable. In short windows, however, sample correlations are noisy, and spurious cross-sector similarities can enter the graph and distort the subspace projection. Restricting the neighbor search to intra-sector assets acts as a structural regularizer on the locality topology, preventing unstable cross-sector connections while preserving the economically meaningful ones. Because the problem it addresses is one of short windows, S-OHRP was designed for, and is evaluated over, windows of up to two years (Section 5.2).

A second, independent motivation concerns the in-sample objective. The OHRP criterion of Eq. (4.7) targets volatility only, which naturally favors stronger denoising re-

gardless of the impact on returns. An investor seeking to maximize risk-adjusted returns has no way to express that preference within OHRP. Generalizing the selection criterion to arbitrary objectives, including maximization, closes this gap, and does so while containing the original OHRP as a special case. The two changes act on different stages, and the ablation of Chapter 5 separates their effects, since the sectoral graph changes the projection itself, while the Sharpe objective changes which projection is selected, at the price of more frequent changes in the selected parameters.

4.2.2 Sector-Constrained Locality Graph

Let $s_n \in \{1, \ldots, S\}$ denote the sector of asset n and $A_s = \{n : s_n = s\}$ the corresponding asset set. The sector-constrained k-nearest neighborhood is defined as $\mathcal{N}_k^s(\boldsymbol{x}_{n'}) = \{\boldsymbol{x}_n : n \in A_{s_{n'}},\ \boldsymbol{x}_n \text{ is among the } k \text{ nearest neighbors of } \boldsymbol{x}_{n'} \text{ within } A_{s_{n'}}\}$, which yields the similarity matrix

$$\tilde{S}_{nn'} = \begin{cases} e^{-\frac{\|\boldsymbol{x}_n - \boldsymbol{x}_{n'}\|^2}{u}} & \text{if } \boldsymbol{x}_n \in \mathcal{N}_k^s(\boldsymbol{x}_{n'}) \text{ or } \boldsymbol{x}_{n'} \in \mathcal{N}_k^s(\boldsymbol{x}_n), \\ 0 & \text{otherwise}, \end{cases} \tag{4.8}$$

where the locality bandwidth $u = \bar{Q}$ is estimated exactly as in Section 4.1.2, and the neighborhood relation is symmetrized in the same way, so that $\tilde{S} = \tilde{S}^T$. When a sector contains fewer than k assets ($|A_s| < k$), all available sector peers are used. The remainder of the projection pipeline, comprising the degree matrix, the Laplacian, and the orthogonal basis construction, carries over unchanged, with $\tilde{S}$ simply replacing S. It is worth observing that defining the sector as the full asset universe recovers the original OHRP similarity matrix, so that S-OHRP is a strict generalization of OHRP rather than a competing design.

4.2.3 Generalized In-Sample Optimization

The in-sample parameter selection of Eq. (4.7) is generalized by jointly optimizing (k, d, r) under any scalar objective F:

$$(k, d, r)_{\text{opt}} = \underset{(k,d,r)}{\operatorname{argmin}}\ \delta \cdot F\left(\boldsymbol{\omega}_{k,d,r} \,|\, \boldsymbol{x}_{\text{in}}\right), \qquad \delta \in \{-1, +1\}, \tag{4.9}$$

where $\delta = +1$ minimizes and $\delta = -1$ maximizes F. Setting $F = \sigma^2_{\text{in}}$ with $\delta = +1$ recovers the OHRP objective. In this work, two S-OHRP instances are evaluated, both with sector-constrained projections:

- **S-OHRP**: $F = \widehat{SR}_{\text{in}}$, $\delta = -1$ (in-sample Sharpe ratio maximization);
- **S-OHRP$_{\text{Vol}}$**: $F = \sigma^2_{\text{in}}$, $\delta = +1$ (in-sample variance minimization).

This pair of instances was not chosen arbitrarily, since comparing S-OHRP$_{\text{Vol}}$ against OHRP isolates the effect of the sectoral constraint, while comparing the two variants against each other isolates the effect of the objective function, thus providing a natural ablation structure for the experiments of Chapter 5.

4.2.4 The Sectoral Gini Index (SGI)

The classical Gini index, applied to portfolio weights with $\sum_i \omega_i = 1$, can be written as

$$GI = \frac{\sum_{i,j} |\omega_i - \omega_j|}{2N}, \tag{4.10}$$

with $GI = 0$ for the equally weighted portfolio and $GI \to (N-1)/N$ under full concentration. Useful as it is, this measure is blind to the sectoral composition of the portfolio, because a portfolio may be well diversified at the asset level and still have its weights clustered within a single economic sector, and holding forty different consumer stocks is, after all, diversification only in a rather narrow sense. To capture this effect, we propose the Sectoral Gini Index (SGI), which applies the same operator to the sector-level aggregate weights $\Omega_s = \sum_{i \in s} \omega_i$:

$$SGI = \frac{\sum_{s,s'} |\Omega_s - \Omega_{s'}|}{2S}, \tag{4.11}$$

capturing cross-sector concentration that the asset-level GI alone cannot detect. The SGI is reported for all strategies in the experiments of Chapter 5 and, as will be seen, it reveals a structural feature of the orthogonal methods that had remained invisible to asset-level concentration measures.

5 EXPERIMENTAL RESULTS AND DISCUSSION

This chapter reports the out-of-sample evaluation of the methods proposed in Chapter 4. Section 5.1 covers OHRP over long estimation windows, from one to five years, where the question is whether the orthogonal projection improves upon HRP and the classical benchmarks when data are relatively abundant. Section 5.2 covers S-OHRP over short windows, from 0.2 to two years, precisely the regime in which estimation noise is most damaging and in which the sectoral prior is expected to pay off. Section 5.3 then discusses the two sets of results jointly, since it is only by reading them together that the division of labor between the two methods becomes visible.

Before proceeding, a methodological remark is in order. The results reported here are not identical to those published in (Pergher; Soldera; Scharcanski, 2026b). The published study employed a universe of 351 Brazilian assets (including delisted stocks) observed from 2008 to 2025, with 21-day rebalancing periods. For this dissertation, all experiments were re-executed under a single unified protocol, with the same universe of 158 assets (defined by the index constituents of February 2026, a look-ahead in the selection of the sample whose implications are discussed in Section 5.1.1 and among the limitations of Chapter 6), the same evaluation period (2011–2026), the same 42-day rebalancing frequency (chosen a priori, for the reasons given in Section 5.1.1), the same transaction cost model, and the same statistical methodology, so that the two contributions of this work can be compared directly and fairly. The resulting figures naturally differ in their absolute values from the published ones, although, as will be seen, every qualitative finding of the original article is preserved, which is in itself reassuring evidence of the robustness of the method. A second remark concerns numerical sensitivity. Because OHRP and S-OHRP select their projection parameters by an in-sample criterion over a discrete grid, floating-point differences of the order of machine precision, such as those produced by a different parallel configuration of the linear-algebra routines, can occasionally change the selected parameters at a given rebalancing date and, through them, the composed Sharpe ratio of an individual window, by a median of about 0.02 and by up to about 0.17. The benchmark strategies are not affected, and Appendix B explains the mechanism and how the reported results can be reproduced exactly. The interpretations that follow therefore rest on patterns that recur across window lengths and on regime averages rather than on any single window.

5.1 Orthogonal Hierarchical Risk Parity: Experimental Results

This section evaluates OHRP over long estimation windows, from one to five years, against the three classical benchmarks. The regime is deliberately the favorable one for correlation-based methods, since abundant in-sample data keep the estimated similarities reliable, and the question is therefore whether the orthogonal projection still adds value where the noise problem it addresses is least severe.

5.1.1 Experimental Setup

The experimental protocol follows the unified configuration used throughout this dissertation. The asset universe comprises 158 Brazilian stocks belonging to the IBOV and Small Caps (SMLL) indices as of February 2026, classified into 10 B3 economic sectors, with daily adjusted closing prices retrieved from Yahoo Finance spanning January 2011 to February 2026. The SELIC rate serves as the risk-free benchmark.

Portfolios are built with a rolling in-sample window of length $WL \in \{1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0\}$ years and held for 42 business days (approximately two months), after which the portfolio is rebalanced, resulting in 89 out-of-sample holding periods per configuration. No forecast is involved, since the weights computed at each rebalancing date are held without trading, drifting only with relative prices (Eq. (2.1)), over the next 42 business days, and the out-of-sample evaluation uses the daily returns realized over that holding period, after which the window rolls forward by 42 days, so that the 89 holding periods do not overlap and, chained together, cover the whole evaluation period. The holding period was fixed before any of the experiments reported here were run, and was not selected by comparing alternatives on these data. The starting point was the published OHRP study (Pergher; Soldera; Scharcanski, 2026b), which used 21 business days, the one-month holding period that is the usual convention in out-of-sample portfolio evaluation (DeMiguel; Garlappi; Uppal, 2009; López de Prado, 2016) and whose suitability for OHRP that study had already established. Doubling it to 42 days was a deliberate trade-off, because it halves the number of rebalancings, and with it the turnover and transaction costs of the adaptive strategies, which re-select their projection parameters at every rebalancing, and it also halves the computational cost of the simulations, since each rebalancing requires a full grid search, keeping the experiments of this dissertation within a budget that allows them to be easily reproduced. Only assets priced on the last

day of the in-sample window and with at least 90% of valid observations within it are considered at each rebalancing date, an observation being valid when the daily return is available and nonzero, since zero returns in the price source typically signal stale prices on days without trading, while the remaining invalid observations of an eligible asset are set to zero, and assets with null in-sample volatility are discarded, so that the number of eligible assets $N(t)$ varies over time, between 19 and 150 across dates and window lengths (median around 90). The universe itself, however, is defined by the index constituents of February 2026 and therefore contains only surviving stocks, a form of look-ahead in the selection of the sample. Because every strategy, the benchmarks included, draws from this same universe, the bias inflates absolute performance levels more than it distorts the comparisons between strategies on which the conclusions of this work rest (Chapter 6). Transaction costs of 3 basis points are applied to the traded volume at the first day of each holding period. OHRP is compared against HRP, Risk Parity (RP) with the linear Ledoit–Wolf shrinkage estimator (Ledoit; Wolf, 2004b), and the Equally Weighted (EW) portfolio. Out-of-sample performance is assessed by the metrics defined in Chapter 2, and all pairwise statistical comparisons are computed over the 89 paired holding-period observations at the 95% confidence level, using the Wilcoxon signed-rank test as the primary criterion and Student's paired t-test as a secondary reference, for the reasons set out in Section 2.6.

5.1.2 Results and Analysis

The results are examined in three groups: first the risk-adjusted return measures, then the risk measures proper, and finally the concentration and trading-activity measures that reveal what the risk reduction costs. This ordering is not neutral, since the three groups tell rather different stories about the same portfolios.

5.1.2.1 Risk-Adjusted Returns

Considering the composed (net-of-costs) Sharpe ratio over the full evaluation period, OHRP achieves the best values among all strategies at eight of the nine window lengths, with its strongest results at $WL = 4.5$ years (SR $= 0.343$), $WL = 3.0$ years (SR $= 0.324$), and $WL = 1.0$ year (SR $= 0.312$). The HRP benchmark, for its part, is remarkably stable across windows, with a composed SR that remains between 0.239

and 0.260 regardless of the window length, while RP and EW stay below HRP at every window, and below OHRP at every window except $WL = 1.5$ years. That window is the single exception to the advantage of OHRP, since there it falls behind HRP (0.121 against 0.239) and even below RP and EW, whereas at the remaining windows its margin over HRP ranges from 0.010 ($WL = 2.0$) to 0.086 ($WL = 4.5$). Averaged over all nine windows, the composed SR of OHRP is 0.276, against 0.250 for HRP. This isolated dip, which reappears at 1.2, 1.4, and 1.8 years in the S-OHRP experiment (Section 5.2.2), echoes the published evaluation (Pergher; Soldera; Scharcanski, 2026b), in which OHRP also failed to lead at isolated windows (1.0 and 4.0 years) and recorded an anomalously low result at 3.5 years. The advantage of OHRP is therefore not uniform across window lengths, and, as will become clear, it is precisely this fragility at individual windows that motivates the sectoral extension of Section 4.2. The Sortino ratio results follow the same qualitative pattern. Figure 5.1 traces the corresponding cumulative return paths at three representative window lengths. At $WL = 2.5$ and 4.5 years OHRP stays ahead of every benchmark through nearly the whole sample, the separation building up gradually rather than resulting from any single favorable sub-period, whereas at $WL = 1.5$ years it leads until early 2016 and falls behind HRP from then on, never to recover.

One nuance deserves an honest mention. Although the composed Sharpe ratios favor OHRP at most window lengths, the per-portfolio Wilcoxon tests indicate no statistically significant superiority of OHRP over any competitor in Sharpe terms, and a single significant win in Sortino terms, at the individual holding-period level (the remaining significant wins in these metrics are those of HRP over RP and EW, and of RP over EW), which reflects just how dispersed short-horizon risk-adjusted returns are in the Brazilian market.

5.1.2.2 Risk Measures

The most robust finding of this evaluation concerns risk. OHRP attains the lowest annualized volatility at every single window length, with mean values between 0.142 and 0.152, against 0.160–0.170 for HRP, 0.174–0.182 for RP, and approximately 0.201 for EW. According to the Wilcoxon tests, this volatility reduction is statistically significant against all three benchmarks at all nine window lengths, which amounts to a complete sweep. The same full dominance holds for maximum drawdown (OHRP mean between -5.4% and -5.8% per holding period, against -6.2% to -6.5% for HRP and -7.8% to -7.9% for EW) and for the Pain Index. These results confirm, now under the stricter

unified protocol, the central claim of (Pergher; Soldera; Scharcanski, 2026b), namely that the combination of an orthogonal locality-preserving projection with in-sample volatility minimization yields consistently less volatile portfolios, with smaller maximum drawdowns and a lower Pain Index, than HRP, RP, and EW, and this advantage extends across the entire range of window lengths from 1 to 5 years. Figure 5.2 makes the contrast visible at the level of the individual portfolios, since the volatility distribution of OHRP is shifted left of every competitor at every window shown, whereas the Sharpe ratio distributions overlap substantially, which is precisely why the Wilcoxon tests separate OHRP from the benchmarks on risk but not on risk-adjusted return.

5.1.2.3 Concentration and Turnover

The improved risk profile does not come for free, and its price is paid in concentration and trading activity. OHRP produces the most concentrated portfolios of the comparison, with an asset-level Gini index between 0.69 and 0.71 across all windows, against 0.34–0.39 for HRP and 0.19–0.23 for RP (EW being perfectly diversified by construction). At the sector level, OHRP also exhibits the highest concentration (sectoral Gini index between 0.44 and 0.48), while RP turns out to be the most sector-diversified strategy at every window length, an observation we will return to in Section 5.2. As for trading activity, OHRP is the most active strategy, with turnover ratios between 0.28 and 0.38 per rebalancing (roughly two and a half times the HRP turnover and five times that of RP and EW), which is a direct consequence of re-selecting the projection parameters at every rebalancing date. Notably, OHRP turnover declines as the window length grows, from 0.382 at $WL = 1.0$ year to around 0.31 between 3.0 and 4.0 years and 0.283 at 5.0 years, though not monotonically (it rises from 0.344 to 0.357 between 1.5 and 2.0 years and to 0.326 at 4.5 years), which suggests that longer in-sample histories stabilize the selected projectors.

Table 5.1 reports the composed-series measures for every strategy and window length, that is, the statistics of the single daily return series obtained by chaining all rebalancing windows, net of transaction costs. It is the source of the Sharpe and Sortino figures quoted in this section, whereas the figures for volatility, drawdown, concentration, and turnover are per-portfolio means. Because a composed series compresses 89 holding periods into a single number, it says nothing about the dispersion behind that number, and the corresponding per-portfolio means, with their 95% confidence intervals, are therefore reported in Table A.1 of Appendix A, and it is there that the overlap between the Sharpe

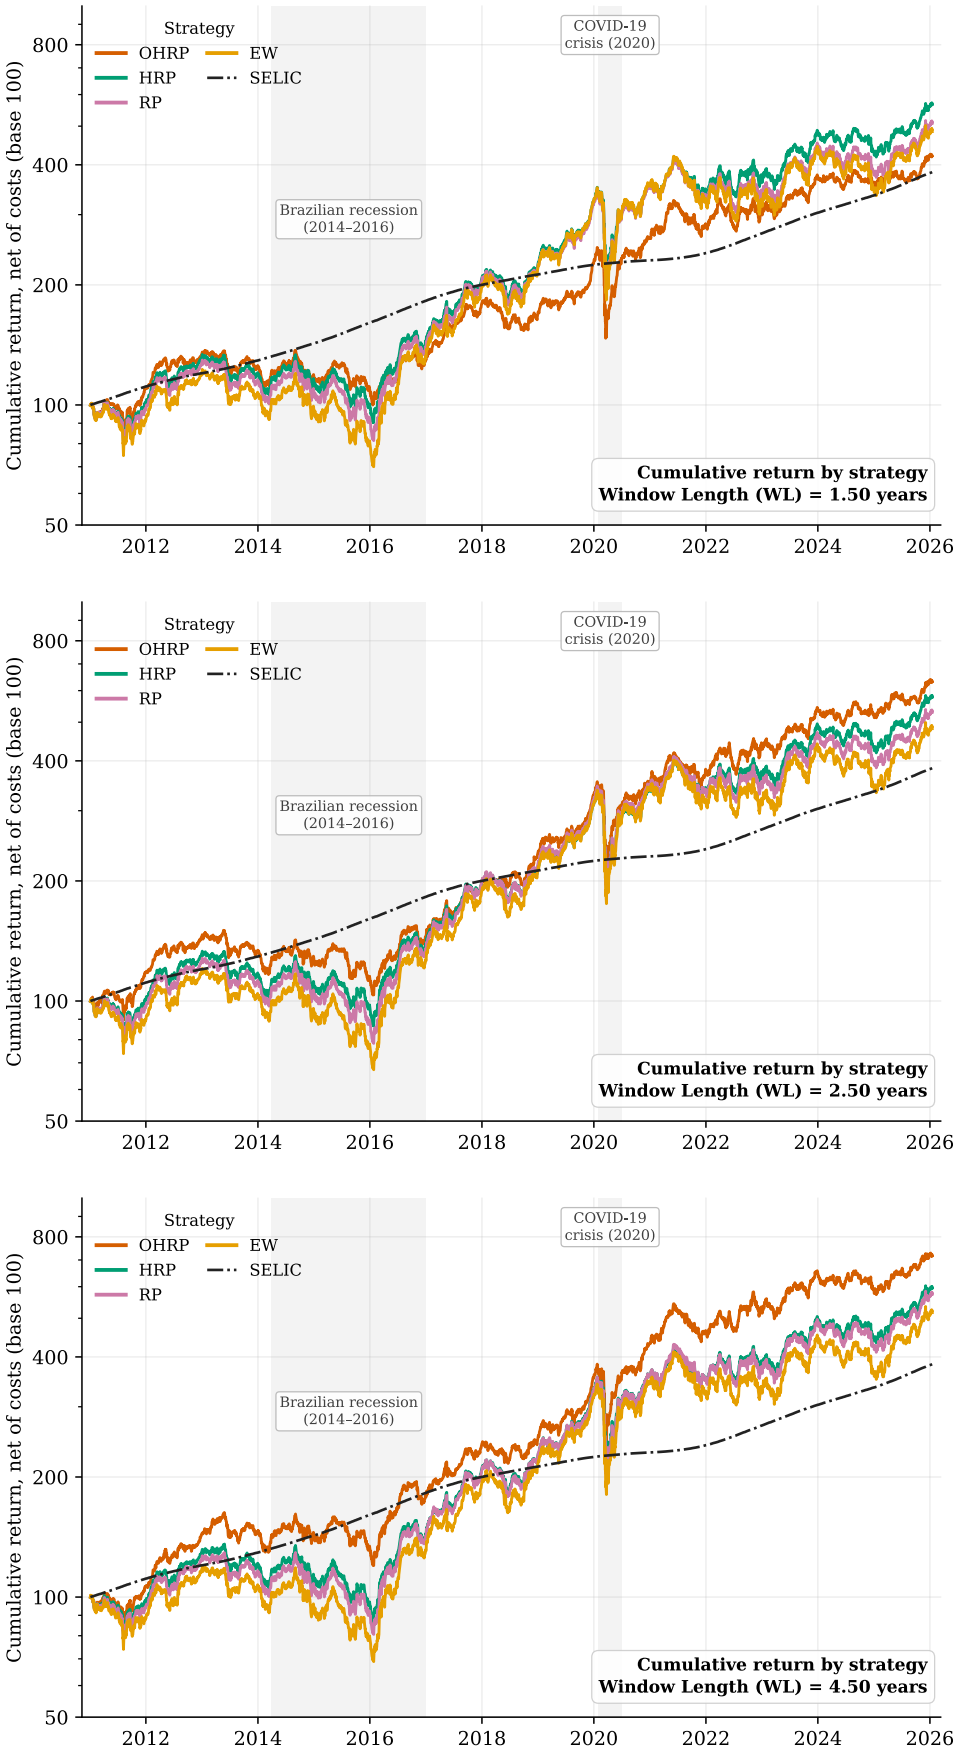


Figure 5.1 – Out-of-sample cumulative returns (net of costs) of OHRP and the benchmark strategies for representative window lengths: $WL = 1.5$ (top), 2.5 (middle), and 4.5 years (bottom).

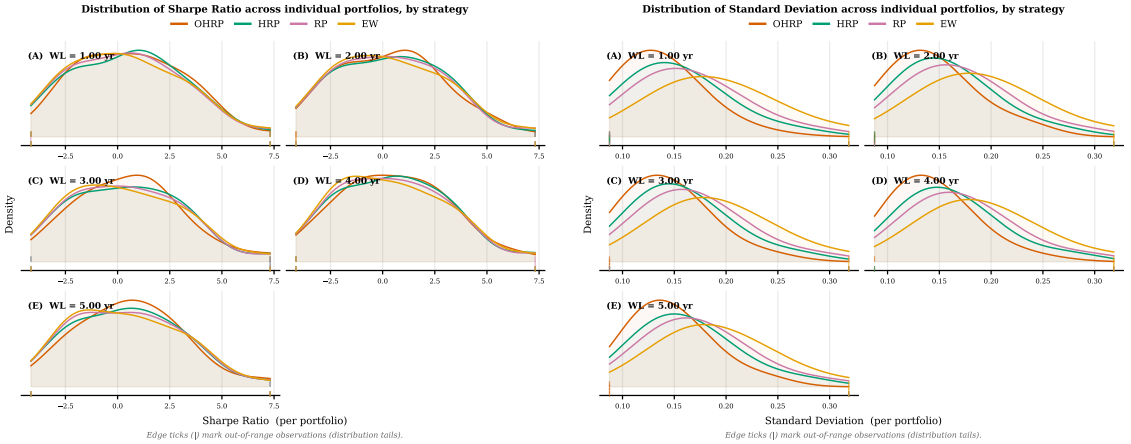


Figure 5.2 – Distributions of per-holding-period Sharpe ratio (left) and annualized volatility (right) across the 89 out-of-sample portfolios, per strategy and window length.

distributions discussed above can be read directly from the interval half-widths.

Table 5.2 consolidates the Wilcoxon signed-rank outcomes. For each metric, the entry counts the number of statistically significant pairwise wins ($p < 0.05$) accumulated by each strategy against the three competitors across the nine window lengths (maximum of 27).

Table 5.1 – Composed-series performance measures for the OHRP experiment, net of transaction costs (3 bps), computed on the full daily return series obtained by chaining all rebalancing windows. Best strategy per window length in bold.

	Estimation window length WL (years)								
Strategy	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
Sharpe ratio									
OHRP	**0.312**	0.121	**0.262**	**0.291**	**0.324**	**0.257**	**0.268**	**0.343**	**0.310**
HRP	0.260	**0.239**	0.252	0.246	0.252	0.241	0.248	0.257	0.253
RP	0.221	0.197	0.207	0.212	0.238	0.222	0.232	0.241	0.236
EW	0.199	0.184	0.180	0.182	0.206	0.197	0.201	0.205	0.211
Sortino ratio									
OHRP	**0.289**	0.111	**0.246**	**0.268**	**0.305**	**0.239**	**0.250**	**0.324**	**0.291**
HRP	0.243	**0.223**	0.237	0.230	0.236	0.227	0.233	0.242	0.238
RP	0.208	0.186	0.195	0.200	0.224	0.209	0.219	0.228	0.223
EW	0.191	0.175	0.172	0.174	0.197	0.188	0.193	0.196	0.202
Annualized mean return									
OHRP	**0.143**	0.113	0.136	**0.141**	**0.147**	0.136	0.137	**0.150**	**0.144**
HRP	0.140	**0.136**	**0.139**	0.138	0.139	**0.138**	0.139	0.141	0.141
RP	0.136	0.131	0.133	0.135	0.140	0.137	**0.139**	0.141	0.140
EW	0.137	0.134	0.133	0.134	0.139	0.137	0.138	0.139	0.140
Annualized volatility									
OHRP	**0.160**	**0.162**	**0.165**	**0.164**	**0.165**	**0.166**	**0.164**	**0.166**	**0.166**
HRP	0.180	0.181	0.181	0.182	0.183	0.185	0.185	0.187	0.188
RP	0.194	0.195	0.196	0.197	0.197	0.198	0.198	0.199	0.201
EW	0.222	0.222	0.222	0.222	0.222	0.222	0.222	0.222	0.223
Maximum drawdown									
OHRP	**-0.406**	**-0.408**	**-0.413**	**-0.410**	**-0.400**	**-0.412**	**-0.390**	**-0.378**	**-0.375**
HRP	-0.427	-0.431	-0.420	-0.434	-0.416	-0.416	-0.407	-0.418	-0.419
RP	-0.451	-0.450	-0.450	-0.455	-0.445	-0.445	-0.443	-0.444	-0.446
EW	-0.481	-0.477	-0.478	-0.480	-0.476	-0.475	-0.473	-0.472	-0.475
Pain Index									
OHRP	**0.060**	**0.069**	**0.070**	**0.069**	**0.062**	**0.062**	**0.065**	**0.064**	**0.066**
HRP	0.078	0.080	0.080	0.084	0.078	0.081	0.082	0.083	0.080
RP	0.091	0.095	0.094	0.096	0.090	0.090	0.091	0.092	0.091
EW	0.114	0.116	0.115	0.117	0.110	0.111	0.110	0.110	0.108

Table 5.2 – Consolidated dominance counts across all nine window lengths (maximum 27 wins per cell). Values are Wilcoxon signed-rank wins, with the paired t-test count given in parentheses where it differs. Arrows indicate the favorable direction of each metric. Column-best in bold.

Strategy	SR↑	SoR↑	$\mu\uparrow$	$\sigma\downarrow$	DD↑	PI↓	GI↓	SGI↓	TO↓
OHRP	0	1 (0)	0	**27**	**27**	**27**	0	0	0
HRP	**11** (13)	**9** (11)	0	18	18	18	9	17	9
RP	2 (7)	2 (4)	0	9	9	9	18	**27**	18
EW	0	0	0	0	0	0	**27**	9	**21**

5.2 Sectoral Orthogonal Hierarchical Risk Parity: Experimental Results

Attention now shifts to the short-window regime, from 0.2 to two years, where estimation noise is most damaging and where the sectoral prior is expected to earn its keep. Two further strategies join the comparison, the Sharpe-driven S-OHRP and its variance-driven counterpart S-OHRP$_{\text{Vol}}$, so that the effect of the sectoral constraint can be separated from the effect of the objective function.

5.2.1 Experimental Setup

The dataset, evaluation period, cost model, and statistical methodology are identical to those of Section 5.1.1: the same 158 Brazilian stocks drawn from the IBOV and SMLL indices, across 10 B3 sectors, from January 2011 to February 2026, 42-day holding periods, 89 out-of-sample portfolios per configuration, 3 basis points of transaction costs, and the same pair of hypothesis tests at the 95% confidence level. The window lengths now cover the short regime, $WL \in \{0.20, 0.40, 0.60, 0.80, 1.00, 1.20, 1.40, 1.60, 1.80, 2.00\}$ years. Six strategies are compared: S-OHRP, S-OHRP$_{\text{Vol}}$, OHRP, HRP, RP, and EW. All grid-searched methods share the parameter grid of Section 4.1.3, and OHRP uses it without sectoral constraints ($\delta = +1$, $F = \sigma^2_{\text{in}}$) and thus serves as the direct ablation baseline, while HRP, RP, and EW require no hyperparameters at all.

The sectoral composition of the universe, depicted in Figure 5.3, is dominated by Consumer Discretionary (41 firms, 25.9%), followed by Financials (21, 13.3%), Industrial Goods (20, 12.7%), Utilities (18, 11.4%), Basic Materials (15, 9.5%), Consumer Staples (14, 8.9%), Health Care (11, 7.0%), Oil, Gas and Biofuels (9, 5.7%), Information Technology (6, 3.8%), and Communications (3, 1.9%).

5.2.2 Results and Analysis

With six strategies and ten window lengths, the results are best read through an organizing principle rather than metric by metric. The one adopted here is the split into two window-length regimes, established first, after which the same three groups of measures used in Section 5.1.2 are examined within that division, closing with the cross-strategy correlations.

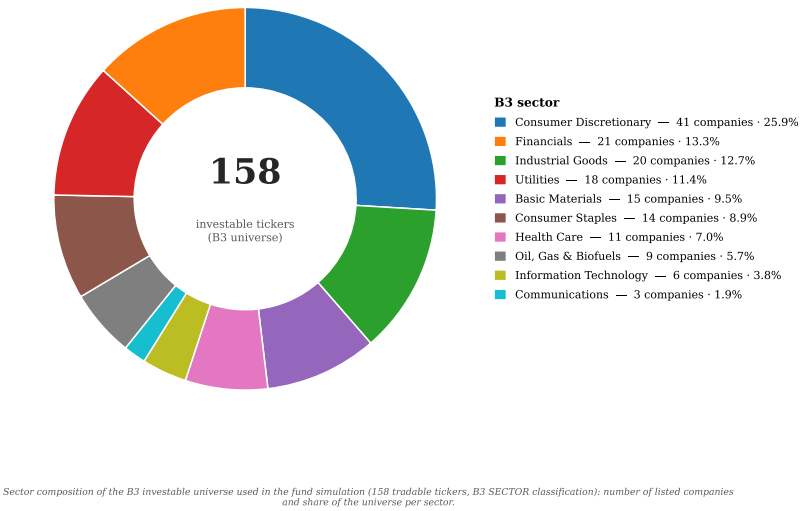


Figure 5.3 – Sectoral composition of the 158-asset universe according to the B3 economic sector classification.

5.2.2.1 Two Window-Length Regimes

Before diving into the individual metrics, it helps to organize the windows into two regimes. In the *short-window regime* ($WL \leq 1.00$ year) estimation noise is at its highest, and in the *long-window regime* ($WL \geq 1.20$ years) in-sample histories are richer. The expectation that motivated the sectoral extension was that the prior would matter most in the first regime and fade in the second, but the results bear out only the first half of that expectation. Averaged over each regime, the composed Sharpe ratio of S-OHRP is 0.356 in the short regime, against 0.305 for OHRP and 0.272 for HRP, and 0.297 in the long regime, against 0.205 and 0.245, which shows that, far from stabilizing, the unconstrained projection weakens at 1.2, 1.4, and 1.8 years, so that the sectoral prior pays off on average in both regimes and most clearly against OHRP in the long one. One possible reason, which these experiments do not test, is that windows of this length usually include at least one episode of market stress, when measured correlations across all pairs rise, largely as an effect of higher volatility (Forbes; Rigobon, 2002), and in such episodes the unconstrained graph admits cross-sector neighbors that do not persist out of sample, and the sectoral prior excludes them. S-OHRP was designed for, and evaluated only up to, two-year windows, so whether its advantage extends to the longer windows of Section 5.1 remains open. Because the result of a single window can move by a few hundredths, and occasionally by more than a tenth, under numerically equivalent executions (see the remark at the opening of this chapter), the regime averages and the patterns that recur across windows are the more reliable basis for interpretation. The

discussion below keeps this division.

5.2.2.2 Risk-Adjusted Returns

S-OHRP achieves the best composed risk-adjusted performance of the entire study at $WL = 1.20$ years, with a net-of-costs Sharpe ratio of 0.482 and Sortino ratio of 0.459, together with the highest composed mean return (0.186). It also leads the composed Sharpe ratio at three of the five short windows, $WL \in \{0.20, 0.60, 1.00\}$ years (0.328, 0.398, and 0.395, against 0.176, 0.370, and 0.312 for OHRP), outperforming HRP and all remaining benchmarks at each of them. The short regime is not uniform, however, since at $WL = 0.40$ year S-OHRP$_{\text{Vol}}$ leads (0.313, against 0.251 for S-OHRP), and at $WL = 0.80$ year OHRP and S-OHRP are virtually tied (0.411 and 0.409). In the long-window regime, beyond its peak at $WL = 1.20$ years, S-OHRP gives way to its variance-driven counterpart, S-OHRP$_{\text{Vol}}$, which attains the best composed Sharpe ratio at $WL \in \{1.40, 1.60, 1.80\}$ years (0.349, 0.296, and 0.271), and OHRP leads at $WL = 2.00$ years (0.262). S-OHRP$_{\text{Vol}}$ surpasses OHRP in Sharpe terms at four of the five long windows, most clearly 0.362 against 0.169 at $WL = 1.20$ years and 0.349 against 0.135 at $WL = 1.40$ years, while OHRP retains the advantage at $WL = 2.00$ years. HRP, meanwhile, is the most consistent strategy in Sharpe and Sortino terms across the full spectrum, accumulating the largest number of statistically significant pairwise wins in these metrics, although it never attains the top composed value at any window.

Taken together, these results support the central hypothesis behind the sectoral extension, though in a broader form than originally framed, since the sectoral prior pays off on average in both the short and the long regime, though not at every window, regularizing the locality topology without destroying signal. Averaged over all ten windows, S-OHRP attains a composed Sharpe ratio of 0.327, against 0.302 for S-OHRP$_{\text{Vol}}$, 0.255 for OHRP, and 0.259 for HRP. At the level of individual holding periods, however, the differences are rarely significant, given that the mean return of S-OHRP is significantly higher than that of OHRP only at $WL \in \{1.20, 1.40\}$ years, and its Sharpe ratio exceeds those of S-OHRP$_{\text{Vol}}$ and HRP only at $WL = 0.80$ year (and those of RP and EW at a few short and intermediate windows), so the composed advantage is an aggregate effect that builds up over the sample rather than a dominance at the level of individual holding periods. Figure 5.4 shows the cumulative return paths of the six strategies at three representative windows. At $WL = 1.20$ years S-OHRP takes the lead in 2012 and ends the sample about 50% above the second-best strategy, whereas at $WL = 0.40$ year it leads for most

of the sample but is overtaken by S-OHRP$_{\text{Vol}}$ in the final years and ends behind HRP as well, which illustrates how uneven the Sharpe-driven selection can be from one window to the next. At $WL = 2.00$ years, finally, the lead alternates among the projection-based methods and HRP, with OHRP and HRP finishing virtually tied and S-OHRP behind both.

5.2.2.3 Risk Measures

On the risk side, the conclusions of Section 5.1 largely carry over to the short windows, with one qualification. OHRP wins 43 of the 50 pairwise Wilcoxon comparisons in volatility, and all seven exceptions are against S-OHRP$_{\text{Vol}}$, which is significantly less volatile than OHRP at five windows ($WL \in \{0.40, 0.60, 0.80, 1.40, 1.80\}$), more volatile at three, and tied at two. Against the three benchmarks the volatility dominance of OHRP is complete, whereas in Pain Index it fails against HRP at $WL \in \{0.20, 1.20\}$ years, and in maximum drawdown at $WL = 0.20$ year. S-OHRP$_{\text{Vol}}$ is thus comparable to OHRP in risk, since, besides being less volatile at five windows and more volatile at three, it is statistically indistinguishable from OHRP in maximum drawdown and in Pain Index at eight of the ten windows, with OHRP ahead at the remaining two, and it also beats HRP in volatility and Pain Index at all ten windows and in maximum drawdown at eight. Since S-OHRP$_{\text{Vol}}$ is significantly less risky than S-OHRP in every risk measure and at every window in the Wilcoxon tests, it is the objective function, rather than the sectoral constraint, that primarily drives the risk profile. S-OHRP itself, whose parameters are selected for in-sample Sharpe rather than variance, exhibits the widest dispersion of per-portfolio returns and volatilities among the projection-based methods at every window, as one would expect.

5.2.2.4 Concentration, Turnover, and the SGI

The asset-level Gini index reproduces the pattern already seen in Section 5.1. The orthogonal methods are again the most concentrated (GI between 0.60 and 0.71), EW is perfectly diversified by construction, and RP and HRP cluster at low concentration values. The novel SGI, however, adds a diagnostic that the asset-level measure cannot provide. RP emerges as the most sector-diversified strategy, winning all pairwise SGI comparisons in both regimes, while OHRP, despite its optimized projection, concentrates its allocations within a few dominant sectors and wins only one SGI comparison per window, always against S-OHRP. S-OHRP is, in fact, the most sector-concentrated strategy

of the study, with mean SGI between 0.54 and 0.59, against 0.47–0.48 for OHRP, and S-OHRP$_{\text{Vol}}$ (0.45–0.48) is statistically indistinguishable from OHRP at nine of the ten windows. The sectoral constraint therefore does not, by itself, diversify across sectors, and combined with Sharpe maximization it concentrates. A plausible explanation, which we state as a hypothesis since the experiments do not isolate it, is that restricting the locality graph to intra-sector edges makes the projected space cluster by sector, so that the HRP tree built on it tends to split along sector lines, while the in-sample Sharpe criterion favors configurations tilted towards the sectors that performed well within the estimation window. Figure 5.5 displays the four distributions side by side and makes the point graphically, since the Gini and Sectoral Gini panels order the strategies differently, confirming that the two concentration measures capture distinct properties of the same portfolios.

Turnover, finally, is the price of Sharpe-driven adaptivity, as S-OHRP is the most active strategy of the study, with turnover ratios between 0.45 and 0.66 per rebalancing, since re-optimizing the parameters for in-sample Sharpe changes the selected projector far more often than variance minimization does. S-OHRP$_{\text{Vol}}$, for its part, trades significantly less than OHRP at all ten windows. EW and RP remain the most passive strategies, followed by HRP.

5.2.2.5 Cross-Strategy Correlations

A final question worth asking is whether S-OHRP is genuinely a different strategy, or merely a refined OHRP. The correlation analysis of the out-of-sample return series suggests the former, since S-OHRP is less correlated with OHRP (between 0.82 and 0.89, depending on the window length) than the benchmark strategies are among themselves, which cluster between 0.98 and 1.00, as Figure 5.6 shows for every window length. The combination of sectoral constraint and Sharpe-driven selection thus produces a strategy that is structurally distinct from its parent, and therefore potentially valuable in multi-strategy combinations, a point we return to in the concluding chapter.

Table 5.3 reports the composed-series measures for all six strategies across the ten window lengths, net of transaction costs, and is the source of the figures quoted throughout this section. Its per-portfolio counterpart is Table A.2 of Appendix A, which adds the two weight-based measures absent from the composed series, the Gini and Sectoral Gini indices, together with turnover, and is therefore the reference for the concentration analysis that follows.

Tables 5.4 and 5.5 consolidate the Wilcoxon signed-rank outcomes for the short-

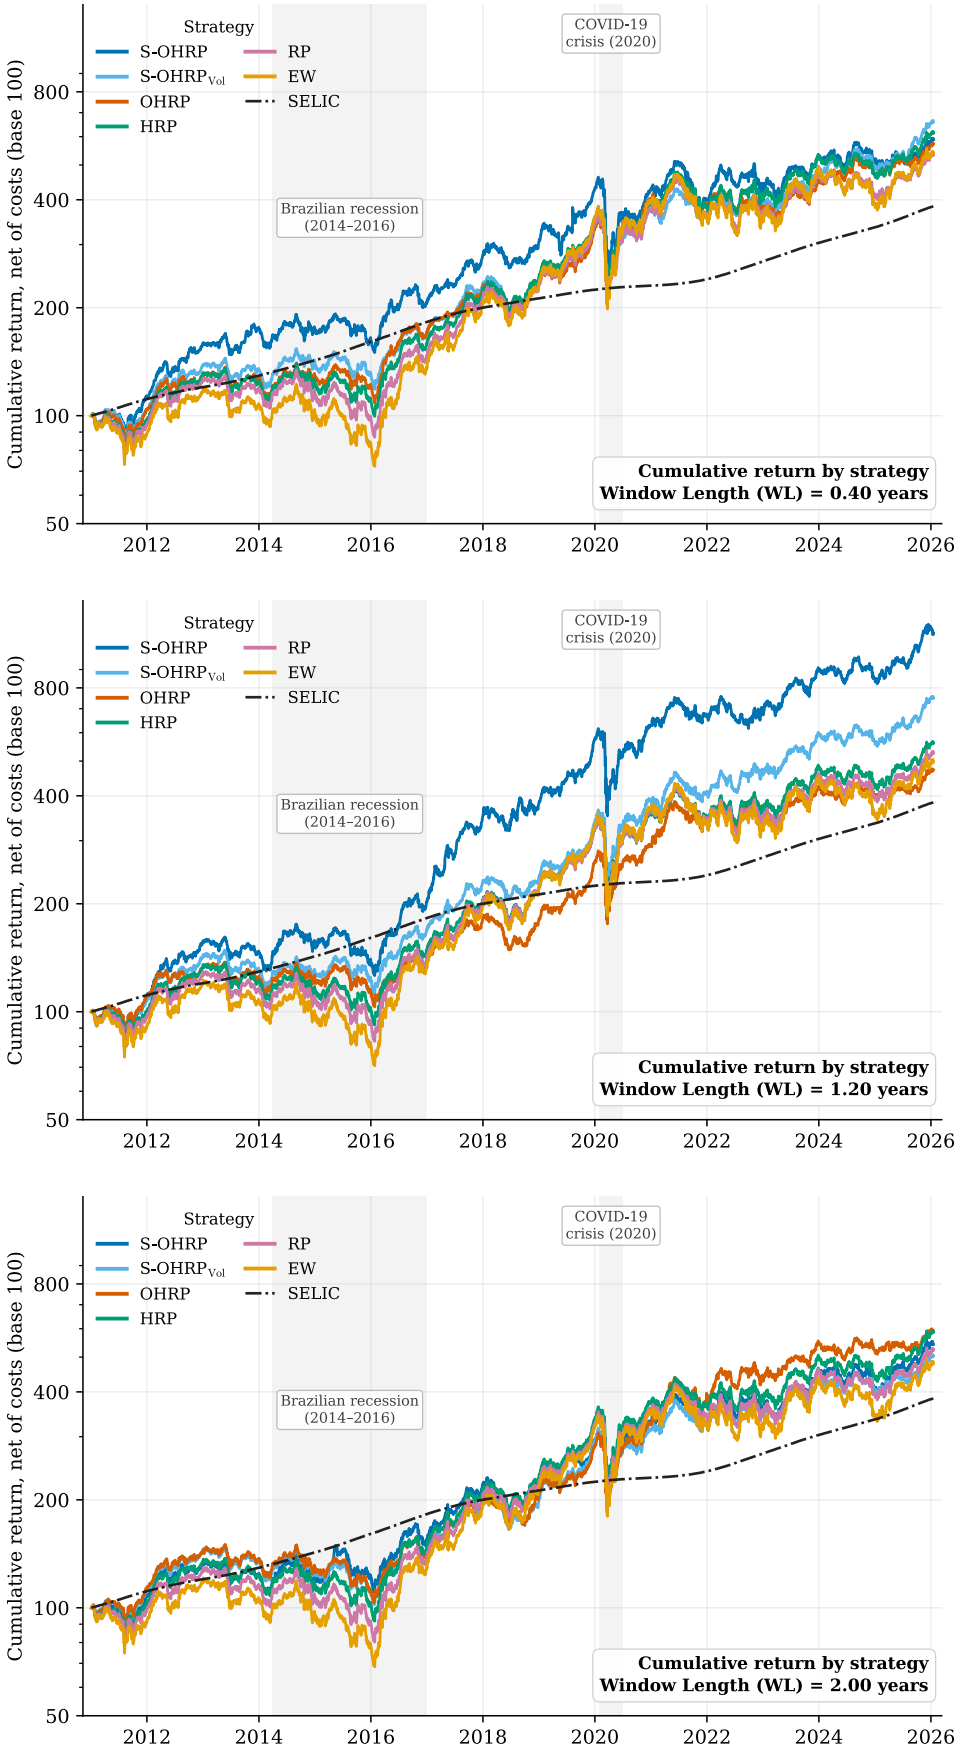


Figure 5.4 – Out-of-sample cumulative returns (net of costs) of the six strategies for representative window lengths: $WL = 0.40$ (top), 1.20 (middle), and 2.00 years (bottom).

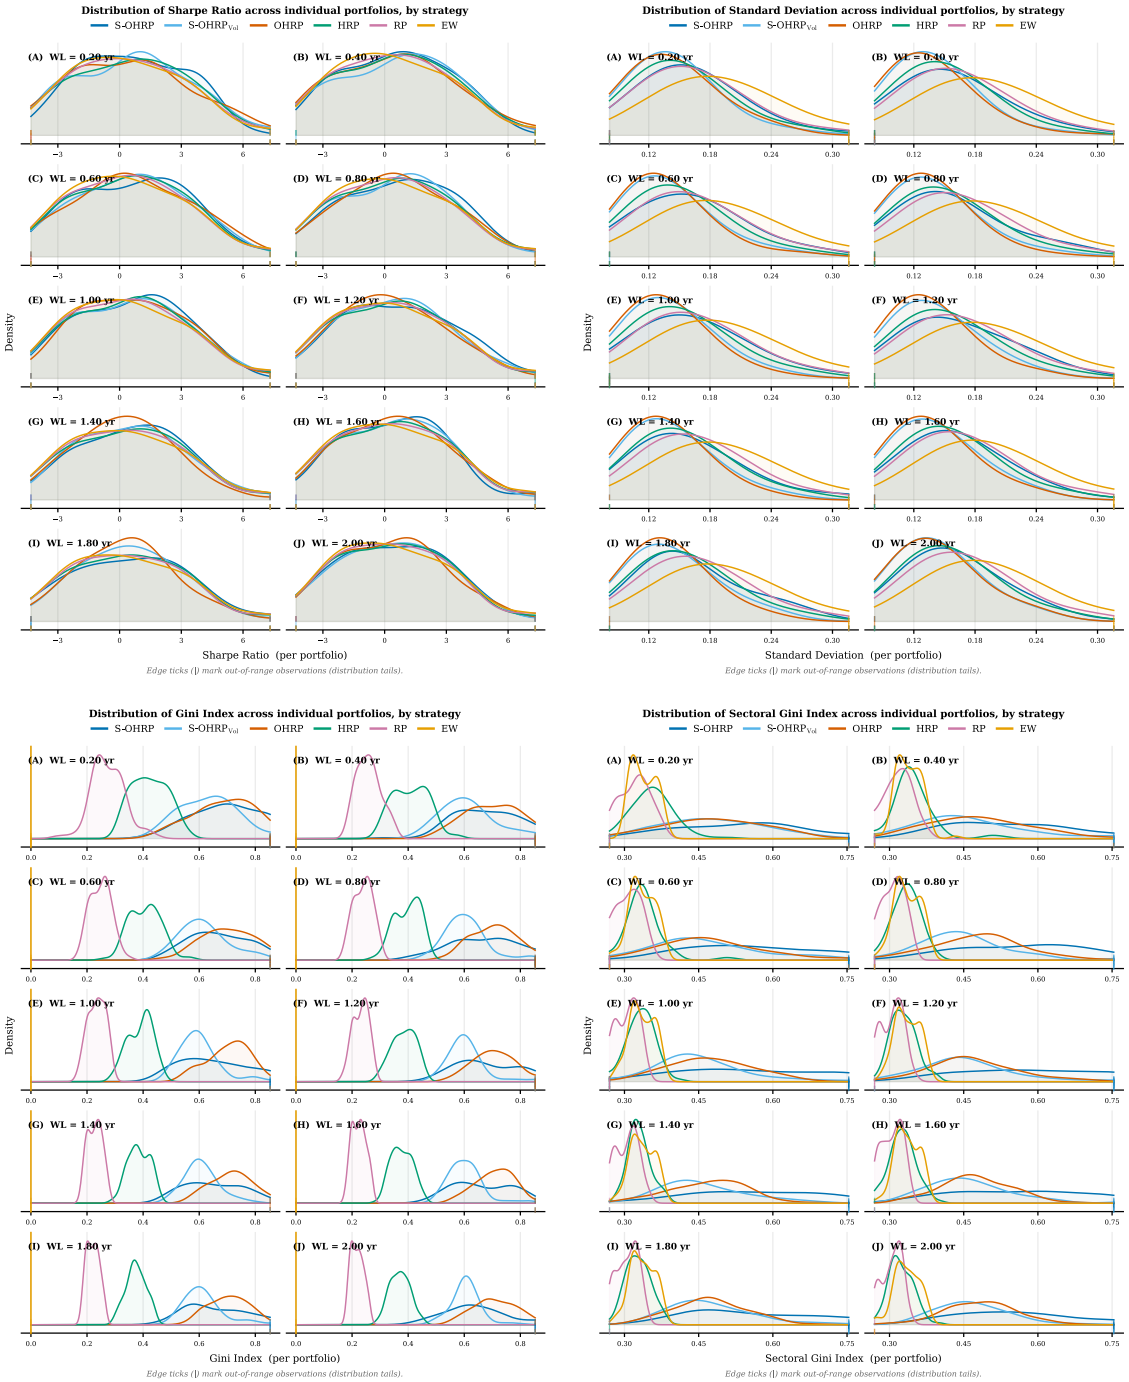


Figure 5.5 – Distributions of per-holding-period Sharpe ratio (top left), annualized volatility (top right), Gini index (bottom left), and Sectoral Gini Index (bottom right) across the 89 out-of-sample portfolios, per strategy and window length.

Table 5.3 – Composed-series performance measures for the S-OHRP experiment, net of transaction costs (3 bps), computed on the full daily return series obtained by chaining all rebalancing windows. Best strategy per window length in bold.

	Estimation window length WL (years)									
Strategy	0.2	0.4	0.6	0.8	1.0	1.2	1.4	1.6	1.8	2.0
Sharpe ratio										
S-OHRP	**0.328**	0.251	**0.398**	0.409	**0.395**	**0.482**	0.304	0.260	0.218	0.221
S-OHRP$_{\text{Vol}}$	0.311	**0.313**	0.385	0.244	0.296	0.362	**0.349**	**0.296**	**0.271**	0.194
OHRP	0.176	0.254	0.370	**0.411**	0.312	0.169	0.135	0.269	0.188	**0.262**
HRP	0.275	0.273	0.283	0.270	0.260	0.236	0.254	0.248	0.234	0.252
RP	0.195	0.215	0.221	0.218	0.221	0.209	0.199	0.212	0.196	0.207
EW	0.212	0.216	0.202	0.209	0.199	0.191	0.182	0.187	0.181	0.180
Sortino ratio										
S-OHRP	**0.303**	0.240	**0.369**	0.376	**0.370**	**0.459**	0.282	0.245	0.206	0.206
S-OHRP$_{\text{Vol}}$	0.289	**0.293**	0.356	0.227	0.276	0.339	**0.325**	**0.273**	**0.252**	0.181
OHRP	0.167	0.238	0.346	**0.382**	0.289	0.157	0.126	0.250	0.176	**0.246**
HRP	0.257	0.257	0.265	0.252	0.243	0.220	0.238	0.232	0.219	0.237
RP	0.182	0.202	0.207	0.204	0.208	0.196	0.188	0.200	0.185	0.195
EW	0.202	0.207	0.193	0.199	0.191	0.182	0.174	0.179	0.173	0.172
Annualized mean return										
S-OHRP	**0.157**	0.142	**0.169**	**0.171**	**0.169**	**0.186**	**0.151**	0.141	0.134	0.133
S-OHRP$_{\text{Vol}}$	0.146	**0.144**	0.156	0.133	0.142	0.153	0.151	**0.142**	**0.138**	0.125
OHRP	0.123	0.134	0.153	0.159	0.143	0.120	0.115	0.137	0.123	0.136
HRP	0.142	0.141	0.143	0.142	0.140	0.136	0.139	0.138	0.136	**0.139**
RP	0.131	0.134	0.136	0.135	0.136	0.134	0.132	0.134	0.131	0.133
EW	0.140	0.141	0.138	0.139	0.137	0.136	0.133	0.135	0.133	0.133
Annualized volatility										
S-OHRP	0.195	0.193	0.190	0.190	0.193	0.194	0.191	0.186	0.186	0.182
S-OHRP$_{\text{Vol}}$	0.170	0.163	0.164	0.165	0.165	0.165	0.165	0.165	0.168	0.167
OHRP	**0.169**	**0.160**	**0.161**	**0.159**	**0.160**	**0.158**	**0.159**	**0.162**	**0.161**	**0.165**
HRP	0.180	0.177	0.178	0.180	0.180	0.181	0.181	0.181	0.181	0.181
RP	0.194	0.192	0.193	0.194	0.194	0.195	0.194	0.195	0.195	0.196
EW	0.222	0.222	0.222	0.222	0.222	0.222	0.222	0.222	0.222	0.222
Maximum drawdown										
S-OHRP	-0.435	-0.483	-0.442	-0.449	-0.428	-0.429	-0.460	**-0.388**	**-0.385**	**-0.404**
S-OHRP$_{\text{Vol}}$	-0.413	-0.392	-0.405	-0.400	**-0.394**	-0.396	**-0.391**	-0.402	-0.399	-0.405
OHRP	**-0.402**	**-0.386**	**-0.401**	**-0.389**	-0.406	**-0.375**	-0.397	-0.410	-0.404	-0.413
HRP	-0.426	-0.425	-0.429	-0.432	-0.427	-0.426	-0.431	-0.424	-0.429	-0.420
RP	-0.460	-0.456	-0.452	-0.451	-0.451	-0.452	-0.452	-0.448	-0.447	-0.450
EW	-0.487	-0.483	-0.481	-0.482	-0.481	-0.480	-0.480	-0.476	-0.476	-0.478
Pain Index										
S-OHRP	0.081	0.078	0.080	0.070	0.068	0.067	0.082	0.088	0.082	0.071
S-OHRP$_{\text{Vol}}$	**0.072**	**0.068**	**0.066**	0.073	0.069	**0.063**	**0.066**	0.067	**0.066**	0.076
OHRP	0.082	0.071	0.078	**0.066**	**0.060**	0.066	0.072	**0.062**	0.070	**0.070**
HRP	0.077	0.079	0.080	0.080	0.078	0.082	0.078	0.079	0.081	0.080
RP	0.092	0.094	0.093	0.092	0.091	0.092	0.093	0.093	0.095	0.094
EW	0.110	0.112	0.115	0.114	0.114	0.116	0.116	0.117	0.116	0.115

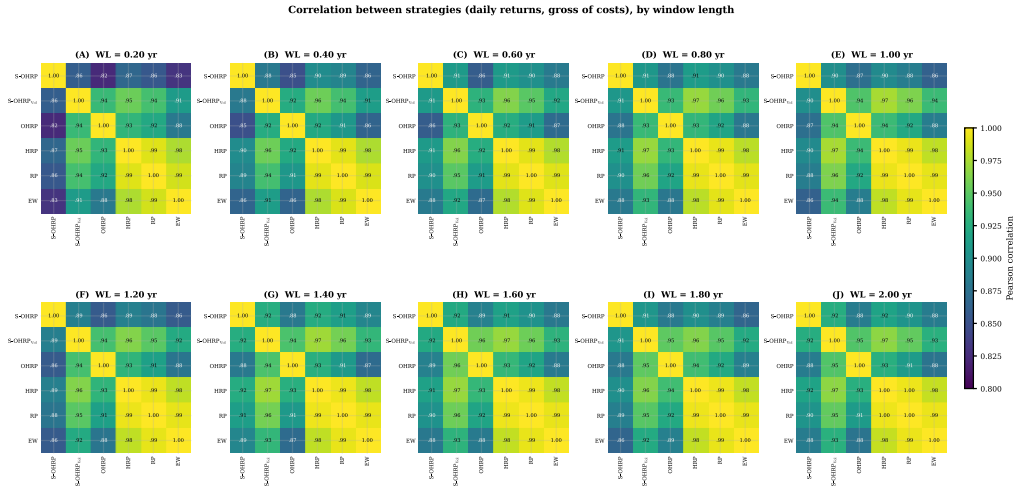


Figure 5.6 – Pairwise correlations between the out-of-sample return series of the six strategies, per window length.

and long-window regimes, respectively. For each metric, the entry counts the number of statistically significant pairwise wins ($p < 0.05$) accumulated by each strategy against the five competitors across the five window lengths of the regime (maximum of 25).

Table 5.4 – Consolidated dominance counts in the short-window regime ($WL \leq 1.00$ year; maximum 25 wins per cell). Values are Wilcoxon signed-rank wins, with the paired t-test count given in parentheses where it differs. Arrows indicate the favorable direction. Column-best in bold.

Strategy	SR↑	SoR↑	$\mu\uparrow$	$\sigma\downarrow$	DD↑	PI↓	GI↓	SGI↓	TO↓
S-OHRP	6 (5)	5 (4)	**4** (1)	8 (5)	8 (7)	7	3	0	0
S-OHRP$_{\text{Vol}}$	2	1 (3)	0	**23** (20)	18 (20)	20	10	6	10
OHRP	0	0	0	21 (23)	**20** (21)	**21** (22)	0	5	5
HRP	**10**	**10**	0	14	14 (11)	12 (11)	15	15	15
RP	0	0	0	5	5	5	20	**25**	20
EW	0	0	0	0	0	0	**25**	16	**25**

Table 5.5 – Consolidated dominance counts in the long-window regime ($WL \geq 1.20$ years; maximum 25 wins per cell). Values are Wilcoxon signed-rank wins, with the paired t-test count given in parentheses where it differs. Arrows indicate the favorable direction. Column-best in bold.

Strategy	SR↑	SoR↑	$\mu\uparrow$	$\sigma\downarrow$	DD↑	PI↓	GI↓	SGI↓	TO↓
S-OHRP	1 (0)	2 (0)	**4** (3)	9	10 (9)	9 (7)	4 (5)	0	0
S-OHRP$_{\text{Vol}}$	7 (6)	5	2 (1)	**22** (20)	20	**20**	10	5	10
OHRP	0	0	0	**22** (24)	**21** (20)	19 (20)	0	5	5
HRP	**9** (10)	**9** (8)	1 (0)	12	10	11 (10)	15	20 (19)	15
RP	0 (1)	0	0	5	5	5	20	**25**	20
EW	0	0	0	0	0	0	**25**	15	**25**

5.3 Discussion

Read together, the two evaluations consolidate the picture drawn in (Pergher; Soldera; Scharcanski, 2026b) and extend it. OHRP is, first and foremost, a risk-reduction device, and its dominance over the three benchmarks in volatility is complete and statistically supported across every window length tested, from 0.2 to 5 years. In maximum drawdown and Pain Index it falls short only against HRP, at $WL = 0.20$ year in both measures and at $WL = 1.20$ years in the Pain Index, and its composed risk-adjusted returns are the best among the compared strategies at eight of the nine long windows. It is worth stressing that the absolute performance levels differ from those originally published, which were obtained with a larger universe including delisted assets, a longer period starting in 2008, and monthly rebalancing, and yet every qualitative conclusion carried over unchanged. In our view, this robustness with respect to the experimental design is one of the more encouraging findings of the re-execution.

The two limitations of OHRP identified in Section 4.2 are visible in the results themselves. The locality graph of Eq. (4.1) is built without any economic prior, so that nearest neighbors are determined purely by the Euclidean distance between return series, and, particularly in shorter windows, spurious cross-sector similarities may contaminate the projection, which is a plausible explanation for the Sharpe ratio dip observed around $WL = 1.5$ years, and at 1.2, 1.4, and 1.8 years in the S-OHRP experiment. The ablation lends direct support to this reading, since S-OHRP$_{\text{Vol}}$, which differs from OHRP only by the sectoral graph, attains a higher composed Sharpe ratio at seven of the ten windows of the short-window experiment, namely $WL \in \{0.20, 0.40, 0.60\}$ years and every long window except 2.00 years. Meanwhile, the in-sample criterion of Eq. (4.7) targets volatility only, which does not necessarily maximize risk-adjusted returns. The sectoral extension addresses both at once, and the experiments show where doing so pays off.

The evidence thus delineates what might be called a division of labor among the proposed methods. For investors whose primary concern is risk, in terms of volatility, maximum drawdown, and persistence of losses (Pain Index), OHRP remains the strategy of choice at any window length, with S-OHRP$_{\text{Vol}}$ as a close alternative at the shorter ones. For investors seeking risk-adjusted returns with estimation windows of up to two years, S-OHRP is the strongest option on average, at the price of the highest turnover and the highest sector concentration of the study, and its advantage stems from the interaction of its two ingredients, since the sectoral prior prevents noisy cross-sector neighbors from

distorting the subspace projection where data are scarce, while the $\delta \cdot F$ framework allows the parameter search to target the Sharpe ratio directly. The advantage is uneven from one window to the next, as $WL \in \{0.40, 1.80, 2.00\}$ years show, where S-OHRP falls below HRP, which is why it should be read as a property of the range of windows rather than as a guarantee at any particular one. At windows of 1.2 years and longer, S-OHRP$_{\text{Vol}}$ nearly matches S-OHRP on average (0.294 against 0.297) with lower risk and turnover, which makes it the natural choice when the sectoral prior is wanted without the costs of Sharpe-driven selection.

A methodological observation is worth extracting from the dominance tables before closing. The Wilcoxon and paired t-test counts, reported side by side in Tables 5.2, 5.4 and 5.5, agree in about three quarters of the cells, and in the OHRP experiment they differ in only five of the thirty-six entries, and every divergence falls in the Sharpe or Sortino columns, with the risk, concentration, and turnover measures matching exactly. The direction of the disagreement is largely systematic, since in four of the five cells it is the t-test that reports *more* significant wins, most visibly for RP, whose Sharpe count rises from two to seven. This is precisely the behavior predicted in Section 2.6. Sharpe and Sortino ratios are ratios of estimated moments, and their sampling distributions over 89 holding periods inherit the heavy tails and skewness of the underlying daily returns far more severely than volatility or maximum drawdown do, so the parametric test, resting on a normality assumption that fails hardest exactly there, over-rejects. The concordance of the two tests on the risk measures, where the OHRP dominance is complete under either procedure, therefore strengthens rather than weakens the central claim of this chapter, while the divergence on risk-adjusted returns is itself an argument for reading those particular comparisons through the distribution-free lens. The per-portfolio tables of Appendix A make the asymmetry concrete. Taking the ratio between the half-width of the 95% confidence interval and the mean itself, the median across all configurations is about 1.1 for the Sharpe ratio and 0.9 for the Sortino ratio, against roughly 0.10 for volatility and 0.16 for maximum drawdown. In other words, the risk-adjusted measures carry an uncertainty comparable to their own magnitude, while the risk measures are estimated an order of magnitude more precisely in relative terms, which is precisely the condition under which a parametric test and a rank-based one may reasonably disagree.

The SGI, for its part, reveals that sector-level concentration is an independent dimension of portfolio structure, genuinely not captured by asset-level measures, as the strategy ranking under SGI (RP first, the orthogonal methods last) differs materially from

the ranking under GI. The sectoral constraint, moreover, does not mitigate the cross-sector concentration of the orthogonal methods, since S-OHRP is the most sector-concentrated strategy of the study, and S-OHRP$_{\text{Vol}}$ is statistically indistinguishable from OHRP at nine of the ten windows. All of the above is subject to one caveat, namely that these observations were obtained in a market with the specific concentration profile of B3, and the range of windows over which the sectoral prior pays off may well differ in more homogeneous markets, a point we revisit in Chapter 6.

6 CONCLUSIONS AND FUTURE WORK

This chapter closes the dissertation. It recapitulates the three contributions and the evidence supporting each, draws out what the results imply for a practitioner allocating capital in the Brazilian market, states plainly the limitations that qualify those implications, points to the directions that seem most worth pursuing next, and finally records the scientific production to which the work gave rise.

This dissertation set out to answer a specific question, namely how to improve the out-of-sample performance of hierarchical risk parity allocation in a large, noisy, sector-concentrated equity universe such as the Brazilian stock market. Three contributions emerged from that effort.

The first is OHRP (Section 4.1), which couples an orthogonal locality-preserving projection, whose parameters (k, d, r) are optimized in-sample at every rebalancing period, to the canonical HRP pipeline. Evaluated over long estimation windows (1 to 5 years) on 158 Brazilian assets from 2011 to 2026, OHRP attained statistically significant lower volatility, smaller maximum drawdowns, and lower Pain Index than HRP, RP, and EW at every single window length, while also achieving the best composed Sharpe ratios at eight of the nine windows. These results were obtained under a stricter, unified protocol, and the fact that they confirm every qualitative finding originally published in (Pergher; Soldera; Scharcanski, 2026b) speaks to the robustness of the method.

The second is S-OHRP (Section 4.2), a strict generalization of OHRP that restricts the locality graph to intra-sector connections and generalizes the in-sample selection into the $\delta \cdot F$ framework, admitting the maximization of arbitrary objectives such as the Sharpe ratio. Evaluated over windows from 0.2 to 2 years, S-OHRP delivered the best risk-adjusted performance of the entire study, a composed Sharpe ratio of 0.482 at $WL = 1.20$ years, and the highest average over the whole range (0.327, against 0.255 for OHRP and 0.259 for HRP), in both the short and the long regime, which, in our reading, indicates that the sectoral prior regularizes the locality topology without destroying signal. The advantage is uneven across individual windows and rarely significant at the level of individual holding periods, and its variance-driven counterpart, S-OHRP$_{\text{Vol}}$, nearly matches it at windows of 1.2 years and longer.

The third is the Sectoral Gini Index (Section 4.2.4), which applies the Gini operator to sector-level aggregate weights. The SGI revealed a structural feature that had remained invisible to asset-level concentration measures, namely that the orthogonal meth-

ods, although diversified across many assets, concentrate their allocations within few sectors, whereas RP is the most sector-diversified strategy of all. S-OHRP, in turn, is the most sector-concentrated strategy of the study, which shows that restricting the graph to sectors does not by itself diversify across them, and this is precisely the kind of risk the SGI was designed to expose.

Beyond the academic contribution, the empirical evidence supports a rather clear decision rule for practitioners in the Brazilian market. Investors whose mandate prioritizes risk control, in the form of lower volatility, smaller maximum drawdowns, and a lower Pain Index, should prefer OHRP, whose dominance in volatility over the benchmarks proved complete across all window lengths tested, from 0.2 to 5 years, and whose advantage in maximum drawdown and Pain Index is nearly as systematic. Investors seeking to maximize risk-adjusted returns with estimation windows of up to two years should prefer S-OHRP, accepting the higher turnover and sector concentration as the price of Sharpe-driven adaptivity. Allocators who want the sectoral prior without these costs may consider S-OHRP$_{\text{Vol}}$, which beats HRP in volatility and Pain Index at every window from 0.2 to 2 years, trades less than OHRP, and nearly matches S-OHRP at windows of 1.2 years and longer. Finally, allocators concerned with sector-level diversification would do well to monitor the SGI alongside asset-level metrics, because the two dimensions of concentration are empirically distinct and, as the OHRP case illustrates, they can point in opposite directions for the very same portfolio.

Seven limitations qualify the conclusions of this work, and it seems only fair to state them plainly. First, the evaluation covers a single market, and the range of windows over which the sectoral prior pays off, as well as the magnitude of the risk-reduction benefits, may differ in markets with different sectoral concentration profiles. Second, the asset universe excludes delisted stocks and, being defined by the index constituents at the end of the sample, amounts to a look-ahead in the selection of assets (a form of data snooping) that introduces survivorship bias. The bias is attenuated, though not eliminated, by the fact that all strategies, the benchmarks included, are compared under the same universe, so that it affects absolute performance levels more than the relative comparisons on which the conclusions rest. Third, transaction costs are modeled as a constant 3 basis points, without market impact or liquidity constraints, which favors high-turnover strategies such as S-OHRP, and the net advantage of Sharpe-driven selection would likely shrink under more punitive cost models. Fourth, the in-sample optimization relies on an exhaustive grid search over a discrete parameter grid, which both bounds the achievable adaptivity

and imposes a nontrivial computational burden at every rebalancing period. Fifth, because the parameters are selected by an in-sample criterion over that discrete grid, the results of the projection-based methods at an individual window are sensitive to floating-point-level differences in the numerical routines, which can move a composed Sharpe ratio by up to about 0.17 (Appendix B), and the conclusions above therefore rest on regime averages and on patterns that recur across windows. Sixth, the paired tests treat the 89 holding periods as independent, an assumption that volatility clustering weakens for the risk measures, so that their p-values are approximate, and dependence-robust tests such as that of Ledoit and Wolf (2008) would sharpen the statistical evidence. Seventh, S-OHRP was designed for the short-window problem and evaluated only up to two-year windows. Its advantage held on average over the longer windows of that range, though not at every one of them, and whether it continues over the longer windows on which OHRP was evaluated is not known.

Several directions for future work follow naturally from this dissertation. The most immediate are two extensions of the evaluation: S-OHRP over the longer windows on which OHRP was evaluated, and both methods in other emerging and developed markets, which would show whether the sectoral prior pays off there as it does here. The second is primarily a data effort rather than a methodological one, since the implementation makes no assumption specific to the Brazilian market and requires only adjusted daily prices, a risk-free rate series, and a sector classification for the target universe, such as GICS. Its computational cost, however, grows with the size of the universe, since every point of the parameter grid entails matrix decompositions whose dimension grows with the number of assets, repeated at each rebalancing. The $\delta \cdot F$ framework, in turn, invites objectives beyond variance and the Sharpe ratio, notably downside-oriented criteria such as the Sortino ratio or CVaR, which seem well suited to the return asymmetries typical of emerging markets. The same framework admits composite objectives, and, since S-OHRP proved the most sector-concentrated strategy of the study, an objective that penalizes the SGI would combine the sectoral prior with diversification across sectors. The locality graph itself could be learned rather than imposed, and replacing the hard sector partition with data-driven or fundamentals-based groupings would generalize S-OHRP beyond exchange classifications. Along the same lines, the Euclidean metric that defines the graph could be replaced by correlation-based, Mahalanobis, or elastic distances, whose relative merits for financial return series remain to be assessed. Appendix B suggests a further change to the selection itself, since normalizing the sign of the eigenvectors, or averaging

the portfolios of the best few grid candidates rather than selecting a single one, would make the selection robust to arbitrary numerical choices and could also reduce the unevenness of the results across windows. The dip of OHRP at 1.2, 1.4, and 1.8 years also deserves direct investigation, for instance by comparing the share of cross-sector neighbors in its graph in windows with and without episodes of market stress. On the allocation side, the low cross-correlations observed between S-OHRP and OHRP (as low as 0.82) suggest that combining them into a single portfolio, for instance through an equal-risk-contribution aggregation such as that of Rubesam (2022), is a promising avenue. Finally, we believe the SGI opens a measurement agenda of its own, as a risk-monitoring and regularization device applicable to any allocation method, not only the ones proposed here.

6.1 Publications Based on this Work

The research developed in this dissertation gave rise to the following scientific production, presented here in chronological order.

6.1.1 Publications in Journals

- PERGHER, K. G. R.; SOLDERA, J.; SCHARCANSKI, J. An Orthogonal Hierarchical Risk Parity Allocation Method for Improved Portfolio Out-of-Sample Performance. **IEEE Access**, v. 14, 2026. DOI: <10.1109/ACCESS.2026.3656702>. Qualis: A1 (current Sucupira classification; ranked A2 at the time of publication).

6.1.2 Publications in Conferences

- PERGHER, K. G. R.; SOLDERA, J.; SCHARCANSKI, J. Dynamic Orthogonal Lower Dimensional Projections for Improving Hierarchical Risk Allocation and Out of Sample Portfolio Returns. In: **2025 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics Companion (CiFer Companion), IEEE Symposium Series on Computational Intelligence (SSCI)**, Trondheim, Norway, 2025. p. 1–5.
  DOI: <10.1109/CiFerCompanion65204.2025.10980404>.

6.1.3 Manuscripts Under Review

- PERGHER, K. G. R.; SOLDERA, J.; SCHARCANSKI, J. Enhancing Orthogonal Hierarchical Risk Parity with Sectoral Information for Out-of-Sample Portfolio Allocation. Manuscript under review, 2026.

REFERENCES

AVELLANEDA, M.; LEE, J.-H. Statistical arbitrage in the US equities market. **Quantitative Finance**, v. 10, n. 7, p. 761–782, 2010.

BAR-JOSEPH, Z.; GIFFORD, D. K.; JAAKKOLA, T. S. Fast optimal leaf ordering for hierarchical clustering. **Bioinformatics**, Oxford University Press, v. 17, n. suppl_1, p. S22–S29, 2001.

BARBOSA FILHO, F. d. H. A crise econômica de 2014/2017. **Estudos Avançados**, Universidade de São Paulo, v. 31, n. 89, p. 51–60, 2017.

BEKAERT, G.; HARVEY, C. R. Emerging markets finance. **Journal of Empirical Finance**, Elsevier, v. 10, n. 1-2, p. 3–55, 2003.

BHOJRAJ, S.; LEE, C. M. C.; OLER, D. K. What's my line? A comparison of industry classification schemes for capital market research. **Journal of Accounting Research**, v. 41, n. 5, p. 745–774, 2003.

BICKEL, P. J.; LEVINA, E. Regularized estimation of large covariance matrices. **The Annals of Statistics**, v. 36, n. 1, p. 199–227, 2008.

BLACK, F.; LITTERMAN, R. Global portfolio optimization. **Financial Analysts Journal**, v. 48, n. 5, p. 28–43, 1992.

BUN, J.; BOUCHAUD, J.-P.; POTTERS, M. Cleaning large correlation matrices: tools from random matrix theory. **Physics Reports**, v. 666, p. 1–109, 2017.

CAI, D. et al. Orthogonal Laplacianfaces for face recognition. **IEEE Transactions on Image Processing**, v. 15, n. 11, p. 3608–3614, 2006.

CALDEIRA, J. F. et al. Portfolio management using realized covariances: Evidence from Brazil. **EconomiA**, Elsevier, v. 18, n. 3, p. 328–343, 2017.

CHAN, L. K. C.; KARCESKI, J.; LAKONISHOK, J. On portfolio optimization: forecasting covariances and choosing the risk model. **The Review of Financial Studies**, v. 12, n. 5, p. 937–974, 1999.

CICIRETTI, V.; PALLOTTA, A. Network risk parity: graph theory-based portfolio construction. **Journal of Asset Management**, v. 25, n. 2, p. 136–146, 2024.

CONNOR, G.; KORAJCZYK, R. A. Performance measurement with the arbitrage pricing theory: a new framework for analysis. **Journal of Financial Economics**, v. 15, n. 3, p. 373–394, 1986.

COSTA, A.; SILVA, C. da; MATOS, P. The Brazilian financial market reaction to COVID-19: A wavelet analysis. **International Review of Economics & Finance**, Elsevier, v. 82, p. 13–29, 2022.

COTTON, P. **Schur Complementary Allocation: A Unification of Hierarchical Risk Parity and Minimum Variance Portfolios**. 2024. ArXiv preprint arXiv:2411.05807. Available from Internet: <https://arxiv.org/abs/2411.05807>.

DEMIGUEL, V.; GARLAPPI, L.; UPPAL, R. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? **The Review of Financial Studies**, v. 22, n. 5, p. 1915–1953, 2009.

ESTRADA, J. Systematic risk in emerging markets: the D-CAPM. **Emerging Markets Review**, Elsevier, v. 3, n. 4, p. 365–379, 2002.

FAMA, E. F.; FRENCH, K. R. Common risk factors in the returns on stocks and bonds. **Journal of Financial Economics**, v. 33, n. 1, p. 3–56, 1993.

FAN, J.; LIAO, Y.; MINCHEVA, M. Large covariance estimation by thresholding principal orthogonal complements. **Journal of the Royal Statistical Society: Series B (Statistical Methodology)**, v. 75, n. 4, p. 603–680, 2013.

FORBES, K. J.; RIGOBON, R. No contagion, only interdependence: measuring stock market comovements. **The Journal of Finance**, v. 57, n. 5, p. 2223–2261, 2002.

GU, S.; KELLY, B.; XIU, D. Empirical asset pricing via machine learning. **The Review of Financial Studies**, v. 33, n. 5, p. 2223–2273, 2020.

GU, S.; KELLY, B.; XIU, D. Autoencoder asset pricing models. **Journal of Econometrics**, v. 222, n. 1, p. 429–450, 2021.

HE, X.; NIYOGI, P. Locality preserving projections. In: **Advances in Neural Information Processing Systems**. [S.l.: s.n.], 2004. v. 16, p. 153–160.

HOBERG, G.; PHILLIPS, G. Text-based network industries and endogenous product differentiation. **Journal of Political Economy**, v. 124, n. 5, p. 1423–1465, 2016.

JAGANNATHAN, R.; MA, T. Risk reduction in large portfolios: why imposing the wrong constraints helps. **The Journal of Finance**, v. 58, n. 4, p. 1651–1683, 2003.

JAIN, P.; JAIN, S. Can machine learning-based portfolios outperform traditional risk-based portfolios? The need to account for covariance misspecification. **Risks**, v. 7, n. 3, p. 74, 2019.

KELLY, B. T.; PRUITT, S.; SU, Y. Characteristics are covariances: a unified model of risk and return. **Journal of Financial Economics**, v. 134, n. 3, p. 501–524, 2019.

LALOUX, L. et al. Noise dressing of financial correlation matrices. **Physical Review Letters**, v. 83, n. 7, p. 1467–1470, 1999.

LEDOIT, O.; WOLF, M. Honey, I shrunk the sample covariance matrix. **The Journal of Portfolio Management**, v. 30, n. 4, p. 110–119, 2004.

LEDOIT, O.; WOLF, M. A well-conditioned estimator for large-dimensional covariance matrices. **Journal of Multivariate Analysis**, v. 88, n. 2, p. 365–411, 2004.

LEDOIT, O.; WOLF, M. Robust performance hypothesis testing with the Sharpe ratio. **Journal of Empirical Finance**, v. 15, n. 5, p. 850–859, 2008.

LEDOIT, O.; WOLF, M. Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets Goldilocks. **The Review of Financial Studies**, v. 30, n. 12, p. 4349–4388, 2017.

LEDOIT, O.; WOLF, M. Analytical nonlinear shrinkage of large-dimensional covariance matrices. **The Annals of Statistics**, v. 48, n. 5, p. 3043–3065, 2020.

LETTAU, M.; PELGER, M. Factors that fit the time series and cross-section of stock returns. **The Review of Financial Studies**, v. 33, n. 5, p. 2274–2325, 2020.

LÓPEZ DE PRADO, M. Building diversified portfolios that outperform out of sample. **The Journal of Portfolio Management**, v. 42, n. 4, p. 59–69, 2016.

MAILLARD, S.; RONCALLI, T.; TEÏLETCHE, J. The properties of equally weighted risk contribution portfolios. **The Journal of Portfolio Management**, v. 36, n. 4, p. 60–70, 2010.

MANTEGNA, R. N. Hierarchical structure in financial markets. **The European Physical Journal B**, v. 11, n. 1, p. 193–197, 1999.

MARCHENKO, V. A.; PASTUR, L. A. Distribution of eigenvalues for some sets of random matrices. **Mathematics of the USSR-Sbornik**, v. 1, n. 4, p. 457–483, 1967.

MARKOWITZ, H. Portfolio selection. **The Journal of Finance**, v. 7, n. 1, p. 77–91, 1952.

MOSKOWITZ, T. J.; GRINBLATT, M. Do industries explain momentum? **The Journal of Finance**, v. 54, n. 4, p. 1249–1290, 1999.

PERGHER, K. G. R.; SOLDERA, J.; SCHARCANSKI, J. Dynamic orthogonal lower dimensional projections for improving hierarchical risk allocation and out of sample portfolio returns. In: **2025 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics Companion (CiFer Companion), IEEE Symposium Series on Computational Intelligence (SSCI)**. Trondheim, Norway: [s.n.], 2025. p. 1–5.

PERGHER, K. G. R.; SOLDERA, J.; SCHARCANSKI, J. Enhancing orthogonal hierarchical risk parity with sectoral information for out-of-sample portfolio allocation. Manuscript under review. 2026.

PERGHER, K. G. R.; SOLDERA, J.; SCHARCANSKI, J. An orthogonal hierarchical risk parity allocation method for improved portfolio out-of-sample performance. **IEEE Access**, v. 14, p. 16885–16903, 2026.

RAFFINOT, T. Hierarchical clustering-based asset allocation. **The Journal of Portfolio Management**, v. 44, n. 2, p. 89–99, 2017.

RAFFINOT, T. **The hierarchical equal risk contribution portfolio**. 2018. SSRN working paper 3237540.

REIS, F. et al. Using hierarchical risk parity in the Brazilian market: An out-of-sample analysis. **Revista Brasileira de Finanças**, v. 21, n. 4, p. 81–103, 2023.

RUBESAM, A. Machine learning portfolios with equal risk contributions: Evidence from the Brazilian market. **Emerging Markets Review**, Elsevier, v. 51, p. 100891, 2022.

RUBESAM, A.; BELTRAME, A. L. Minimum variance portfolios in the Brazilian equity market. **Revista Brasileira de Finanças**, v. 11, n. 1, p. 81–118, 2013.

SALAS-MOLINA, F.; NIN, J. Fast hierarchical risk parity methods for portfolio selection. **Annals of Operations Research**, 2026.

SALAS-MOLINA, F. et al. An empirical evaluation of distance metrics in hierarchical risk parity methods for asset allocation. **Computational Economics**, v. 66, n. 6, p. 5189–5206, 2025.

SCHWENDNER, P. et al. Adaptive seriational risk parity and other extensions for heuristic portfolio construction using machine learning and graph theory. **The Journal of Financial Data Science**, v. 3, n. 4, p. 65–83, 2021.

SHARPE, W. F. Mutual fund performance. **The Journal of Business**, v. 39, n. 1, pt. 2, p. 119–138, 1966.

SOLDERA, J.; BEHAINE, C. A. R.; SCHARCANSKI, J. Customized orthogonal locality preserving projections with soft-margin maximization for face recognition. **IEEE Transactions on Instrumentation and Measurement**, IEEE, v. 64, n. 9, p. 2417–2426, 2015.

SORTINO, F. A.; PRICE, L. N. Performance measurement in a downside risk framework. **The Journal of Investing**, v. 3, n. 3, p. 59–64, 1994.

SOUZA, P. O. de et al. Risk parity in the Brazilian market. **Economics Bulletin**, v. 37, n. 3, p. 1555–1566, 2017.

WARD, J. H. Hierarchical grouping to optimize an objective function. **Journal of the American Statistical Association**, Taylor & Francis, v. 58, n. 301, p. 236–244, 1963.

WILCOXON, F. Individual comparisons by ranking methods. **Biometrics Bulletin**, v. 1, n. 6, p. 80–83, 1945.

APPENDIX A — DETAILED EXPERIMENTAL RESULTS

This appendix reports the individual-portfolio statistics underlying the analyses of Chapter 5. Whereas the composed-series tables of that chapter describe the single daily return series obtained by chaining all rebalancing windows, the tables below characterize the distribution across the 89 individual out-of-sample portfolios of each configuration, reporting the sample mean of each measure together with the half-width of its 95% confidence interval. They are the descriptive counterpart to the Wilcoxon dominance counts of Chapter 5, in that the latter establish which differences are statistically significant, while the former show the magnitudes involved. Both tables are rotated by ninety degrees on the page, so as to accommodate the nine measures side by side while keeping every page of the document in the same A4 portrait format.

A.1 OHRP Experiment: Individual-Portfolio Measures

Table A.1 – Individual-portfolio measures for the OHRP experiment: mean ± 95% confidence interval (Student-t half-width, $n = 89$ portfolios). Best strategy per window length and measure in bold. SR = Sharpe ratio; SoR = Sortino ratio; μ = annualized mean return; σ = annualized standard deviation; DD = maximum drawdown (higher is less negative, hence better); PI = Pain Index; GI = Gini index; SGI = Sectoral Gini Index; TO = turnover ratio.

Strategy	SR	SoR	μ	σ	DD	PI	GI	SGI	TO
WL = 1.00 yr									
OHRP	**0.65 ± 0.58**	**0.80 ± 0.61**	**0.14 ± 0.08**	**0.14 ± 0.02**	**-0.05 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.02	0.47 ± 0.02	0.38 ± 0.02
HRP	0.57 ± 0.59	0.75 ± 0.63	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.39 ± 0.01	0.33 ± 0.01	0.15 ± 0.01
RP	0.48 ± 0.59	0.64 ± 0.62	0.14 ± 0.11	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.23 ± 0.01	**0.30 ± 0.00**	0.08 ± 0.01
EW	0.40 ± 0.59	0.59 ± 0.62	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.07 ± 0.00**
WL = 1.50 yr									
OHRP	0.35 ± 0.57	0.54 ± 0.60	0.11 ± 0.08	**0.14 ± 0.02**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.71 ± 0.01	0.47 ± 0.01	0.34 ± 0.02
HRP	**0.51 ± 0.59**	**0.68 ± 0.61**	**0.14 ± 0.10**	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.38 ± 0.01	0.33 ± 0.01	0.14 ± 0.00
RP	0.43 ± 0.59	0.59 ± 0.61	0.13 ± 0.10	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.22 ± 0.01	**0.31 ± 0.00**	0.07 ± 0.00
EW	0.37 ± 0.58	0.55 ± 0.60	0.13 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.00	**0.06 ± 0.00**
WL = 2.00 yr									
OHRP	0.50 ± 0.59	0.68 ± 0.63	0.14 ± 0.09	**0.15 ± 0.01**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.71 ± 0.02	0.48 ± 0.02	0.36 ± 0.02
HRP	**0.54 ± 0.60**	**0.73 ± 0.64**	**0.14 ± 0.10**	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.37 ± 0.01	0.32 ± 0.01	0.14 ± 0.00
RP	0.44 ± 0.58	0.59 ± 0.60	0.13 ± 0.10	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.22 ± 0.00	**0.31 ± 0.00**	0.06 ± 0.00
EW	0.36 ± 0.58	0.53 ± 0.60	0.13 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.33 ± 0.01	**0.06 ± 0.00**

Continued on next page

Individual-portfolio measures, OHRP experiment – continued from previous page

Strategy	SR	SoR	μ	σ	DD	PI	GI	SGI	TO
WL = 2.50 yr									
OHRP	**0.57 ± 0.57**	**0.76 ± 0.60**	**0.14 ± 0.08**	**0.15 ± 0.02**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.01	0.48 ± 0.01	0.32 ± 0.02
HRP	0.51 ± 0.58	0.68 ± 0.61	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.37 ± 0.01	0.32 ± 0.00	0.13 ± 0.00
RP	0.44 ± 0.58	0.60 ± 0.60	0.13 ± 0.11	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.21 ± 0.00	**0.30 ± 0.00**	0.06 ± 0.00
EW	0.36 ± 0.58	0.52 ± 0.59	0.13 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.33 ± 0.01	**0.06 ± 0.00**
WL = 3.00 yr									
OHRP	**0.62 ± 0.58**	**0.79 ± 0.60**	**0.15 ± 0.08**	**0.15 ± 0.01**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.01	0.46 ± 0.01	0.31 ± 0.02
HRP	0.52 ± 0.58	0.69 ± 0.61	0.14 ± 0.10	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.36 ± 0.01	0.33 ± 0.01	0.13 ± 0.00
RP	0.48 ± 0.58	0.65 ± 0.61	0.14 ± 0.10	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.21 ± 0.00	**0.30 ± 0.00**	0.06 ± 0.00
EW	0.40 ± 0.58	0.56 ± 0.60	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.33 ± 0.01	**0.06 ± 0.00**
WL = 3.50 yr									
OHRP	0.52 ± 0.54	0.67 ± 0.55	0.14 ± 0.08	**0.15 ± 0.02**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.01	0.46 ± 0.01	0.31 ± 0.02
HRP	**0.52 ± 0.59**	**0.71 ± 0.62**	**0.14 ± 0.10**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.35 ± 0.01	0.32 ± 0.00	0.13 ± 0.00
RP	0.46 ± 0.58	0.63 ± 0.60	0.14 ± 0.11	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.21 ± 0.00	**0.31 ± 0.00**	0.06 ± 0.00
EW	0.38 ± 0.58	0.54 ± 0.60	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.33 ± 0.01	**0.06 ± 0.00**
WL = 4.00 yr									
OHRP	**0.55 ± 0.56**	**0.70 ± 0.57**	0.14 ± 0.08	**0.15 ± 0.01**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.01	0.46 ± 0.01	0.31 ± 0.02
HRP	0.52 ± 0.57	0.69 ± 0.58	**0.14 ± 0.10**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.35 ± 0.01	0.32 ± 0.00	0.13 ± 0.00
RP	0.47 ± 0.57	0.63 ± 0.59	0.14 ± 0.11	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.20 ± 0.00	**0.30 ± 0.00**	**0.06 ± 0.00**

Continued on next page

Individual-portfolio measures, OHRP experiment – continued from previous page

Strategy	SR	SoR	μ	σ	DD	PI	GI	SGI	TO
EW	0.39 ± 0.57	0.54 ± 0.59	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	0.06 ± 0.00
WL = 4.50 yr									
OHRP	**0.67 ± 0.58**	**0.93 ± 0.65**	**0.15 ± 0.08**	**0.15 ± 0.01**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.01	0.45 ± 0.01	0.33 ± 0.02
HRP	0.54 ± 0.57	0.71 ± 0.60	0.14 ± 0.10	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.34 ± 0.01	0.32 ± 0.00	0.12 ± 0.00
RP	0.48 ± 0.57	0.64 ± 0.59	0.14 ± 0.11	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.20 ± 0.00	**0.30 ± 0.00**	**0.06 ± 0.00**
EW	0.40 ± 0.57	0.56 ± 0.59	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	0.06 ± 0.00
WL = 5.00 yr									
OHRP	**0.64 ± 0.55**	**0.85 ± 0.59**	**0.15 ± 0.08**	**0.15 ± 0.01**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.69 ± 0.01	0.44 ± 0.01	0.28 ± 0.01
HRP	0.52 ± 0.56	0.69 ± 0.59	0.14 ± 0.10	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.34 ± 0.01	0.32 ± 0.00	0.12 ± 0.01
RP	0.47 ± 0.57	0.61 ± 0.58	0.14 ± 0.11	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.19 ± 0.01	**0.30 ± 0.00**	0.06 ± 0.00
EW	0.39 ± 0.57	0.54 ± 0.58	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.06 ± 0.00**

A.2 S-OHRP Experiment: Individual-Portfolio Measures

Table A.2 – Individual-portfolio measures for the S-OHRP experiment: mean ± 95% confidence interval (Student-t half-width, $n = 89$ portfolios). Best strategy per window length and measure in bold. SR = Sharpe ratio; SoR = Sortino ratio; μ = annualized mean return; σ = annualized standard deviation; DD = maximum drawdown (higher is less negative, hence better); PI = Pain Index; GI = Gini index; SGI = Sectoral Gini Index; TO = turnover ratio.

Strategy	SR	SoR	μ	σ	DD	PI	GI	SGI	TO
WL = 0.20 yr									
S-OHRP	0.63 ± 0.53	0.77 ± 0.53	**0.16 ± 0.10**	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.70 ± 0.03	0.54 ± 0.03	0.66 ± 0.03
S-OHRP$_{\text{Vol}}$	**0.67 ± 0.59**	**0.92 ± 0.64**	0.15 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.64 ± 0.02	0.48 ± 0.02	0.53 ± 0.02
OHRP	0.50 ± 0.62	0.79 ± 0.68	0.12 ± 0.10	**0.15 ± 0.02**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.69 ± 0.02	0.48 ± 0.02	0.58 ± 0.02
HRP	0.65 ± 0.61	0.89 ± 0.66	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.42 ± 0.01	0.35 ± 0.01	0.32 ± 0.02
RP	0.53 ± 0.61	0.75 ± 0.67	0.13 ± 0.11	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.27 ± 0.01	**0.32 ± 0.01**	0.24 ± 0.02
EW	0.46 ± 0.60	0.67 ± 0.64	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.13 ± 0.03**
WL = 0.40 yr									
S-OHRP	0.66 ± 0.60	0.89 ± 0.64	0.14 ± 0.10	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.68 ± 0.03	0.55 ± 0.03	0.58 ± 0.03
S-OHRP$_{\text{Vol}}$	**0.68 ± 0.64**	**0.97 ± 0.73**	**0.15 ± 0.09**	0.14 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.62 ± 0.02	0.47 ± 0.03	0.41 ± 0.02
OHRP	0.60 ± 0.63	0.84 ± 0.69	0.14 ± 0.09	**0.14 ± 0.02**	**-0.06 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.02	0.48 ± 0.02	0.47 ± 0.02
HRP	0.63 ± 0.61	0.85 ± 0.67	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.41 ± 0.01	0.34 ± 0.01	0.22 ± 0.01
RP	0.52 ± 0.60	0.72 ± 0.65	0.13 ± 0.11	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.26 ± 0.01	**0.31 ± 0.01**	0.14 ± 0.01
EW	0.44 ± 0.60	0.64 ± 0.63	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.09 ± 0.01**
WL = 0.60 yr									
S-OHRP	**0.76 ± 0.60**	0.94 ± 0.62	**0.17 ± 0.10**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.68 ± 0.02	0.56 ± 0.03	0.55 ± 0.03

Continued on next page

Individual-portfolio measures, S-OHRP experiment – continued from previous page

Strategy	SR	SoR	μ	σ	DD	PI	GI	SGI	TO
S-OHRP$_{\text{Vol}}$	0.76 ± 0.62	**0.99 ± 0.72**	0.16 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.62 ± 0.02	0.47 ± 0.02	0.36 ± 0.02
OHRP	0.72 ± 0.61	0.94 ± 0.66	0.15 ± 0.09	**0.14 ± 0.02**	**-0.05 ± 0.01**	**0.03 ± 0.01**	0.71 ± 0.02	0.47 ± 0.02	0.43 ± 0.03
HRP	0.63 ± 0.60	0.84 ± 0.65	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.41 ± 0.01	0.34 ± 0.01	0.18 ± 0.01
RP	0.50 ± 0.60	0.68 ± 0.63	0.14 ± 0.11	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.25 ± 0.01	**0.31 ± 0.01**	0.10 ± 0.01
EW	0.42 ± 0.60	0.61 ± 0.62	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.07 ± 0.00**
WL = 0.80 yr									
S-OHRP	**0.82 ± 0.59**	**1.01 ± 0.62**	**0.17 ± 0.10**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.68 ± 0.03	0.57 ± 0.03	0.54 ± 0.03
S-OHRP$_{\text{Vol}}$	0.60 ± 0.60	0.76 ± 0.61	0.13 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.61 ± 0.02	0.45 ± 0.02	0.33 ± 0.02
OHRP	0.74 ± 0.61	0.88 ± 0.62	0.16 ± 0.09	**0.14 ± 0.02**	**-0.05 ± 0.01**	**0.02 ± 0.01**	0.71 ± 0.02	0.48 ± 0.02	0.39 ± 0.02
HRP	0.59 ± 0.60	0.78 ± 0.63	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.40 ± 0.01	0.33 ± 0.01	0.17 ± 0.01
RP	0.49 ± 0.59	0.66 ± 0.61	0.14 ± 0.11	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.24 ± 0.01	**0.31 ± 0.01**	0.08 ± 0.01
EW	0.42 ± 0.59	0.60 ± 0.62	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.07 ± 0.00**
WL = 1.00 yr									
S-OHRP	**0.66 ± 0.57**	**0.86 ± 0.60**	**0.17 ± 0.09**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.65 ± 0.03	0.55 ± 0.03	0.49 ± 0.03
S-OHRP$_{\text{Vol}}$	0.64 ± 0.61	0.85 ± 0.66	0.14 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.60 ± 0.02	0.45 ± 0.02	0.32 ± 0.02
OHRP	0.65 ± 0.58	0.80 ± 0.61	0.14 ± 0.08	**0.14 ± 0.02**	**-0.05 ± 0.01**	**0.03 ± 0.01**	0.70 ± 0.02	0.47 ± 0.02	0.38 ± 0.02
HRP	0.57 ± 0.59	0.75 ± 0.63	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.39 ± 0.01	0.33 ± 0.01	0.15 ± 0.01
RP	0.48 ± 0.59	0.64 ± 0.62	0.14 ± 0.11	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.23 ± 0.01	**0.30 ± 0.00**	0.08 ± 0.01
EW	0.40 ± 0.59	0.59 ± 0.62	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.07 ± 0.00**
WL = 1.20 yr									

Continued on next page

Individual-portfolio measures, S-OHRP experiment – continued from previous page

Strategy	SR	SoR	μ	σ	DD	PI	GI	SGI	TO
S-OHRP	0.67 ± 0.58	0.87 ± 0.60	**0.19 ± 0.10**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.68 ± 0.02	0.59 ± 0.03	0.51 ± 0.04
S-OHRP$_{\text{Vol}}$	**0.69 ± 0.61**	**0.89 ± 0.65**	0.15 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	**0.03 ± 0.01**	0.61 ± 0.02	0.46 ± 0.02	0.30 ± 0.02
OHRP	0.46 ± 0.60	0.69 ± 0.68	0.12 ± 0.08	**0.14 ± 0.01**	**-0.06 ± 0.01**	0.03 ± 0.01	0.71 ± 0.02	0.47 ± 0.02	0.36 ± 0.02
HRP	0.53 ± 0.60	0.72 ± 0.64	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.38 ± 0.01	0.33 ± 0.01	0.15 ± 0.00
RP	0.45 ± 0.59	0.62 ± 0.61	0.13 ± 0.10	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.23 ± 0.01	**0.31 ± 0.00**	0.07 ± 0.00
EW	0.39 ± 0.59	0.57 ± 0.61	0.14 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.06 ± 0.00**
WL = 1.40 yr									
S-OHRP	0.62 ± 0.57	0.80 ± 0.58	**0.15 ± 0.10**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.65 ± 0.02	0.56 ± 0.03	0.46 ± 0.03
S-OHRP$_{\text{Vol}}$	**0.64 ± 0.59**	**0.83 ± 0.61**	0.15 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.61 ± 0.01	0.46 ± 0.02	0.31 ± 0.02
OHRP	0.38 ± 0.56	0.55 ± 0.59	0.12 ± 0.08	**0.14 ± 0.01**	**-0.05 ± 0.01**	**0.03 ± 0.01**	0.71 ± 0.02	0.47 ± 0.02	0.36 ± 0.02
HRP	0.54 ± 0.59	0.71 ± 0.62	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.38 ± 0.01	0.33 ± 0.01	0.14 ± 0.00
RP	0.43 ± 0.58	0.58 ± 0.60	0.13 ± 0.10	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.23 ± 0.01	**0.31 ± 0.00**	0.07 ± 0.00
EW	0.37 ± 0.58	0.54 ± 0.60	0.13 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.06 ± 0.00**
WL = 1.60 yr									
S-OHRP	0.51 ± 0.57	0.69 ± 0.60	**0.14 ± 0.09**	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.66 ± 0.02	0.56 ± 0.03	0.46 ± 0.03
S-OHRP$_{\text{Vol}}$	**0.59 ± 0.58**	0.71 ± 0.60	0.14 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	**0.03 ± 0.01**	0.60 ± 0.01	0.46 ± 0.02	0.29 ± 0.02
OHRP	0.54 ± 0.57	**0.72 ± 0.61**	0.14 ± 0.08	**0.15 ± 0.02**	**-0.06 ± 0.01**	0.03 ± 0.01	0.71 ± 0.01	0.47 ± 0.01	0.36 ± 0.02
HRP	0.51 ± 0.58	0.68 ± 0.60	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.38 ± 0.01	0.33 ± 0.01	0.14 ± 0.00
RP	0.44 ± 0.59	0.61 ± 0.61	0.13 ± 0.10	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.22 ± 0.01	**0.31 ± 0.00**	0.07 ± 0.00
EW	0.37 ± 0.58	0.55 ± 0.60	0.13 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.00	**0.06 ± 0.00**

Continued on next page

Individual-portfolio measures, S-OHRP experiment – continued from previous page

Strategy	SR	SoR	μ	σ	DD	PI	GI	SGI	TO
WL = 1.80 yr									
S-OHRP	0.50 ± 0.62	0.74 ± 0.63	0.13 ± 0.10	0.17 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.66 ± 0.02	0.56 ± 0.03	0.45 ± 0.04
S-OHRP$_{\text{Vol}}$	**0.59 ± 0.56**	**0.80 ± 0.64**	**0.14 ± 0.08**	0.15 ± 0.02	-0.06 ± 0.01	**0.03 ± 0.01**	0.62 ± 0.02	0.48 ± 0.02	0.30 ± 0.02
OHRP	0.49 ± 0.56	0.70 ± 0.62	0.12 ± 0.08	**0.15 ± 0.01**	**-0.06 ± 0.01**	0.03 ± 0.01	0.71 ± 0.02	0.47 ± 0.02	0.36 ± 0.02
HRP	0.52 ± 0.59	0.68 ± 0.61	0.14 ± 0.10	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.37 ± 0.01	0.33 ± 0.01	0.14 ± 0.00
RP	0.43 ± 0.58	0.58 ± 0.60	0.13 ± 0.10	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.22 ± 0.00	**0.31 ± 0.00**	0.07 ± 0.00
EW	0.37 ± 0.58	0.53 ± 0.60	0.13 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.34 ± 0.01	**0.06 ± 0.00**
WL = 2.00 yr									
S-OHRP	0.51 ± 0.59	0.71 ± 0.62	0.13 ± 0.10	0.17 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.65 ± 0.02	0.56 ± 0.03	0.45 ± 0.03
S-OHRP$_{\text{Vol}}$	0.50 ± 0.61	0.68 ± 0.66	0.13 ± 0.09	0.15 ± 0.02	-0.06 ± 0.01	**0.03 ± 0.01**	0.61 ± 0.01	0.47 ± 0.02	0.28 ± 0.02
OHRP	0.50 ± 0.59	0.68 ± 0.63	0.14 ± 0.09	**0.15 ± 0.01**	**-0.06 ± 0.01**	0.03 ± 0.01	0.71 ± 0.02	0.48 ± 0.02	0.36 ± 0.02
HRP	**0.54 ± 0.60**	**0.73 ± 0.64**	**0.14 ± 0.10**	0.16 ± 0.02	-0.06 ± 0.01	0.03 ± 0.01	0.37 ± 0.01	0.32 ± 0.01	0.14 ± 0.00
RP	0.44 ± 0.58	0.59 ± 0.60	0.13 ± 0.10	0.18 ± 0.02	-0.07 ± 0.01	0.03 ± 0.01	0.22 ± 0.00	**0.31 ± 0.00**	0.06 ± 0.00
EW	0.36 ± 0.58	0.53 ± 0.60	0.13 ± 0.12	0.20 ± 0.02	-0.08 ± 0.01	0.04 ± 0.01	**0.00 ± 0.00**	0.33 ± 0.01	**0.06 ± 0.00**

APPENDIX B — REPRODUCIBILITY OF THE EXPERIMENTS

The results of Chapter 5 can be reproduced exactly provided that four elements are matched: the code, the data, the software environment, and one numerical setting that is easy to overlook, the number of threads used by the linear-algebra library. This appendix lists what is needed and explains why the last two matter.

B.1 Code, Data, and Protocol

The implementation of OHRP, S-OHRP, the benchmark strategies, and the SGI, together with the backtest driver and the scripts that produce every table and figure of Chapter 5, is publicly available at <https://github.com/KgPergher98/OHRP>, in the release that accompanies this version of the dissertation. The data are included in the same repository and are all freely available: daily adjusted closing prices retrieved from Yahoo Finance, the effective SELIC rate published by the Central Bank of Brazil, and the B3 economic sector classification of the 158 assets.

The script `simulationFund.py` runs both experiments under the protocol of Section 5.1.1: first rebalancing on 3 January 2011, holding periods of 42 business days, the parameter grid $k \in \{1, 3, \ldots, 13\}$, $d \in \{5, 10, \ldots, 35\}$, $r \in \{0.92, 0.94, \ldots, 1.00\}$ (245 projections per method and rebalancing date), the 90% valid-observation filter, and transaction costs of 3 basis points. The script `analise_experimentos.py`, run once per experiment, turns the stored return series into the tables and figures. The grid search dominates the cost, and reproducing both experiments takes on the order of 80 CPU-hours on the machine below, a time that parallelizes almost linearly across rebalancing dates.

B.2 Reference Environment

All results reported in this dissertation were produced in the environment of Table B.1, whose package versions are pinned in the repository.

Table B.1 – Reference environment of the experiments.

Component	Version
Operating system	Ubuntu 24.04.2 LTS, Linux 6.8.0, x86-64, glibc 2.39
Processor	Intel Core i7-14700
Python	3.11.13
Packages	numpy 2.3.0, scipy 1.16.1, pandas 2.3.0, scikit-learn 1.7.2, Riskfolio-Lib 7.0.1, cvxpy 1.7.2
BLAS/LAPACK	OpenBLAS 0.3.29 (numpy) and 0.3.28 (scipy), `Haswell` kernel
BLAS threads	1 (see Section B.3)
Pinned versions	`requirements-dissertation.txt`

B.3 Why the Number of Threads Matters

OHRP and S-OHRP select their projection parameters (k, d, r) by comparing in-sample objectives that often differ by a fraction of a percent among the best candidates. The selection is therefore sensitive to perturbations at the level of machine precision, and such perturbations arise whenever the order of floating-point operations changes. A multithreaded BLAS splits sums and matrix products among threads and combines the partial results in an order that depends on the number of threads, and the eigensolvers used by the projection inherit these rounding differences, which can flip the sign of an eigenvector (an arbitrary choice from the mathematical point of view) or, in S-OHRP, select a different basis of an eigenspace with repeated eigenvalues, since its sector-restricted graph has one connected component per sector. HRP clusters the projected coordinates by Pearson correlation, which is not invariant to such sign changes, so that a different (k, d, r) can win the grid search, and from that rebalancing date onwards the whole portfolio path differs.

The effect was measured by running the same code on the same machine with all 28 hardware threads and with a single thread. HRP, RP, and EW agreed to within 4×10^{-7}, whereas the composed Sharpe ratio of OHRP and S-OHRP at an individual window changed by a median of about 0.02 and by up to about 0.08 for OHRP and 0.17 for S-OHRP, although, averaged over the windows of each experiment, the changes largely cancel out. The same mechanism explains why results differ across operating systems and across versions of the numerical libraries, even with identical code.

To reproduce the reported figures exactly, fix the number of threads at one before Python starts, by setting `OMP_NUM_THREADS`, `OPENBLAS_NUM_THREADS`, `MKL_NUM_THREADS`, and `NUMEXPR_NUM_THREADS` to 1. The computation is then deterministic, as portfolios of OHRP, S-OHRP, and S-OHRP$_{\text{Vol}}$ recomputed from scratch at dates spread over the whole sample gave identical parameters and weights.

APPENDIX C — RESUMO EXPANDIDO EM PORTUGUÊS

Título: Métodos Ortogonais de Paridade de Risco Hierárquica para Alocação de Carteiras no Mercado de Ações Brasileiro.

A alocação de capital entre ativos de risco é um dos problemas centrais das finanças quantitativas. Desde o arcabouço de média-variância de Markowitz, sabe-se que a solução ótima depende de estimativas da distribuição conjunta dos retornos, extraídas na prática de amostras limitadas e ruidosas, e a otimização amplifica esses erros em alocações instáveis e concentradas, a ponto de a carteira ingênua de pesos iguais (EW) permanecer difícil de superar fora da amostra. O *Hierarchical Risk Parity* (HRP), proposto por López de Prado em 2016, substituiu a inversão da matriz de covariância por três etapas, a saber, agrupamento hierárquico, quasi-diagonalização e bisecção recursiva. É robusto por construção, mas seus insumos continuam sendo correlações amostrais que, em universos grandes e janelas curtas, são fortemente contaminadas por ruído, degradando a hierarquia sobre a qual toda a alocação repousa. O problema agrava-se em mercados emergentes, com volatilidade mais alta e assimétrica e índices concentrados em poucos setores. O mercado brasileiro (B3) entre 2011 e 2026 é um laboratório exigente, abrangendo a recessão de 2014–2016, o rali seguinte, o choque da COVID-19 em 2020 e ciclos da SELIC entre mínimas históricas e dois dígitos. A pergunta respondida aqui é como melhorar o desempenho fora da amostra da alocação hierárquica nesse tipo de universo.

A resposta desenvolve-se em três contribuições. A primeira é o **OHRP** (*Orthogonal Hierarchical Risk Parity*) (Pergher; Soldera; Scharcanski, 2026b), que projeta as séries de retornos em um subespaço ortogonal de menor dimensão preservador de localidade antes das etapas do HRP. A projeção compõe um estágio de Análise de Componentes Principais, que retém uma fração r da variabilidade, com Projeções Ortogonais com Preservação de Localidade (OLPP), construídas a partir de um grafo de k vizinhos mais próximos e do respectivo Laplaciano. Os parâmetros (k, d, r) não são fixados a priori, e sim reotimizados a cada rebalanceamento por busca em grade, minimizando a volatilidade dentro da amostra, o que converte um pré-processamento estático em componente adaptativo da estratégia e fornece ao agrupamento insumos mais estáveis. Essa adaptação não é aprendizado supervisionado, pois não há variável-alvo a prever nem modelo ajustado que se transporte de uma janela para a seguinte, pois a projeção é não supervisionada e a busca em grade apenas escolhe, entre projeções candidatas, a que melhor atende ao objetivo dentro da amostra, refazendo tudo a cada rebalanceamento.

A segunda contribuição, o **S-OHRP** (*Sectoral OHRP*) (Pergher; Soldera; Scharcanski, 2026a), generaliza o OHRP em duas direções. O grafo de localidade passa a admitir apenas conexões intrassetoriais, injetando um *prior* econômico na representação aprendida e barrando vizinhanças espúrias entre setores, particularmente danosas em janelas curtas. E a seleção de parâmetros é unificada em um arcabouço $\delta \cdot F$, no qual δ define a direção da busca e F é qualquer objetivo escalar, o que torna admissíveis critérios como o índice de Sharpe. Avaliam-se duas instâncias, o S-OHRP, que maximiza o Sharpe, e o S-OHRP$_{\text{Vol}}$, que minimiza variância, e, como remover a restrição setorial recupera exatamente o OHRP, a generalização é estrita e o par forma uma ablação natural. A terceira contribuição é o **Índice de Gini Setorial** (SGI), que aplica o operador de Gini aos pesos agregados por setor, $\Omega_s = \sum_{i \in s} \omega_i$, capturando a concentração que o Gini por ativo não detecta, já que uma carteira pode ser diversificada entre muitos papéis e ainda ter seus pesos aglomerados em um único setor.

Os experimentos seguem um protocolo unificado. O universo compreende 158 ações brasileiras dos índices IBOV e Small Caps (SMLL), em 10 setores da B3, com preços diários ajustados de janeiro de 2011 a fevereiro de 2026 e a SELIC como taxa livre de risco. As carteiras são mantidas por 42 dias úteis, gerando 89 períodos fora da amostra por configuração, com custos de 3 pontos-base sobre o volume negociado. Esse período foi fixado antes dos experimentos e dobra os 21 dias úteis do estudo original, que equivalem a um mês, a convenção usual na avaliação fora da amostra (DeMiguel; Garlappi; Uppal, 2009; López de Prado, 2016), para reduzir o *turnover*, os custos de transação e o custo computacional das simulações. O OHRP é avaliado em janelas de 1,0 a 5,0 anos contra HRP, paridade de risco (RP, com o estimador linear de Ledoit–Wolf) e EW, e o S-OHRP, em janelas de 0,20 a 2,00 anos, contra os mesmos *benchmarks* e o próprio OHRP. Nove métricas são medidas, e as comparações par a par adotam o teste de Wilcoxon como critério primário, ao nível de 95%, com o teste t pareado reportado como referência secundária. Embora o estudo original (Pergher; Soldera; Scharcanski, 2026b) tenha empregado 351 ativos, incluindo deslistados, de 2008 a 2025 com rebalanceamento de 21 dias, todos os experimentos foram reexecutados sob o protocolo acima, e suas constatações qualitativas foram preservadas.

Em janelas longas, o OHRP alcançou a menor volatilidade anualizada nas nove janelas testadas, com médias entre 0,142 e 0,152, contra 0,160–0,170 do HRP, 0,174–0,182 da RP e cerca de 0,201 do EW, com significância completa nos testes de Wilcoxon, 27 vitórias em 27 comparações, e dominância igualmente completa em *drawdown* máx-

imo e *Pain Index*. Obteve ainda os melhores índices de Sharpe líquidos de custos em oito das nove janelas, com picos de 0,343 em 4,5 anos e 0,324 em 3,0 anos e média de 0,276 contra 0,250 do HRP. A única exceção é a janela de 1,5 ano (0,121), queda que reaparece em 1,2, 1,4 e 1,8 ano no experimento do S-OHRP, de modo que, como no estudo original, a vantagem não é uniforme entre janelas, fragilidade que motivou a extensão setorial. O preço dessa melhoria é a concentração, com Gini entre 0,69 e 0,71, e o *turnover* cerca de duas vezes e meia o do HRP, decorrente da reseleção dinâmica dos parâmetros.

No experimento de janelas curtas e intermediárias (0,2 a 2 anos), o S-OHRP entregou o melhor desempenho ajustado ao risco de todo o estudo em $WL = 1{,}20$ ano, com Sharpe composto de 0,482 e Sortino de 0,459, e liderou também em 0,20, 0,60 e 1,00 ano (0,328, 0,398 e 0,395, contra 0,176, 0,370 e 0,312 do OHRP). Na média das dez janelas alcançou 0,327, contra 0,302 do S-OHRP$_{\text{Vol}}$, 0,255 do OHRP e 0,259 do HRP, e a vantagem aparece nos dois regimes, com 0,356 nas janelas curtas e 0,297 nas longas, nas quais o OHRP sem restrição enfraquece (0,205). A vantagem é irregular entre janelas e raramente significativa carteira a carteira, e o S-OHRP$_{\text{Vol}}$ lidera em 0,40 e de 1,40 a 1,80 ano e quase empata com o S-OHRP na média das janelas longas (0,294 contra 0,297). O conjunto sustenta a hipótese central numa forma mais ampla do que a prevista, já que o *prior* setorial rende, em média, nos dois regimes, embora não em todas as janelas. Como uma janela isolada pode variar alguns centésimos, e às vezes mais de um décimo, entre execuções numericamente equivalentes, a leitura apoia-se nas médias por regime.

O SGI revelou uma característica invisível às medidas clássicas, a de que os métodos ortogonais, embora diversificados entre muitos ativos, concentram suas alocações em poucos setores, ao passo que a RP é a mais diversificada setorialmente, com 25 vitórias em 25 comparações por regime. A restrição setorial não mitiga essa concentração, pois o S-OHRP é a estratégia mais concentrada por setor de todo o estudo (SGI médio entre 0,54 e 0,59, contra 0,47–0,48 do OHRP), e o S-OHRP$_{\text{Vol}}$ é estatisticamente indistinguível do OHRP em quase todas as janelas. Uma explicação plausível, ainda não isolada pelos experimentos, é que o grafo intrassetorial faz o espaço projetado agrupar-se por setor, de modo que a árvore do HRP tende a dividir-se segundo os setores, enquanto o critério de Sharpe favorece configurações inclinadas aos setores que foram bem na janela de estimação. O S-OHRP mostrou-se ainda menos correlacionado com o OHRP, entre 0,82 e 0,89, do que os *benchmarks* entre si, agrupados entre 0,98 e 1,00, confirmando tratar-se de estratégia estruturalmente distinta. Do ponto de vista metodológico, os dois testes concordam na quase totalidade das comparações e, no experimento do OHRP, onde divergem

é nas métricas de retorno ajustado, com o teste t apontando mais vitórias na maioria dos casos, comportamento esperado quando a normalidade falha e que reforça a escolha não paramétrica. Os dois testes, porém, supõem diferenças independentes entre períodos, o que vale bem para as medidas de retorno, mas é enfraquecido pelo agrupamento de volatilidade nas medidas de risco, cujos p-valores devem ser lidos como aproximados.

A evidência delineia uma regra de decisão clara. Investidores que priorizam controle de risco devem preferir o OHRP, cuja dominância em volatilidade sobre os *benchmarks* é completa em todas as janelas testadas, de 0,2 a 5 anos, e cuja vantagem em *drawdown* máximo e *Pain Index* é quase tão sistemática. Já quem busca maximizar retorno ajustado ao risco com janelas de estimação de até dois anos deve preferir o S-OHRP, aceitando o maior *turnover* e a maior concentração setorial como preço da adaptividade, e quem quer o *prior* setorial sem esses custos pode optar pelo S-OHRP$_{\text{Vol}}$, que vence o HRP em volatilidade e *Pain Index* em todas as janelas de 0,2 a 2,0 anos, negocia menos que o OHRP e quase iguala o S-OHRP nas janelas de 1,2 ano ou mais. Por fim, alocadores atentos à diversificação setorial devem monitorar o SGI ao lado das métricas por ativo, já que as duas dimensões são empiricamente distintas. Sete limitações as qualificam: a avaliação cobre um único mercado; o universo exclui deslistados e, por ser definido pelos componentes ao fim da amostra, equivale a usar informação futura na seleção dos ativos, com viés de sobrevivência atenuado por comparar todas as estratégias, *benchmarks* inclusive, sob o mesmo universo, o que afeta mais os níveis absolutos do que as comparações; os custos são constantes e sem impacto de mercado, o que favorece estratégias de alto *turnover*; a busca em grade discreta limita a adaptividade e impõe custo computacional; e, por selecionar parâmetros nessa grade, os métodos de projeção são sensíveis a diferenças numéricas da ordem da precisão de máquina, que podem mover o Sharpe composto de uma janela em até cerca de 0,17; os testes pareados supõem independência entre os períodos, de modo que os p-valores das medidas de risco são aproximados; e o S-OHRP, desenhado para janelas curtas, só foi avaliado até dois anos. Entre os trabalhos futuros estão a extensão do S-OHRP a janelas longas e a outros mercados, que exige sobretudo dados, já que a implementação não depende do mercado brasileiro, embora o custo computacional cresça com o número de ativos; objetivos de risco de cauda como Sortino e CVaR, ou com penalização do SGI, no arcabouço $\delta \cdot F$; o aprendizado do grafo de localidade a partir de dados, inclusive com métricas alternativas à distância euclidiana; a combinação das estratégias pouco correlacionadas em uma única carteira; e o uso do SGI como instrumento de monitoramento.

APPENDIX D — DECLARAÇÃO DE USO DE INTELIGÊNCIA ARTIFICIAL GENERATIVA

Na elaboração desta dissertação, o autor utilizou a ferramenta de inteligência artificial generativa Claude (Anthropic, modelo Fable 5) como apoio à redação, estruturação e revisão do texto, à tradução de conteúdo e à formatação do documento LaTeX. Todo o conteúdo científico, incluindo os métodos propostos, a implementação computacional, os experimentos, os resultados e as conclusões, é de autoria integral do autor, tendo sido previamente desenvolvido e publicado nos artigos que fundamentam este trabalho. Todas as saídas da ferramenta foram revisadas, verificadas e validadas pelo autor, que assume integral responsabilidade pelo conteúdo final, nos termos da Portaria CNPq nº 2.664/2026.