



Index Tracking Portfolio Construction with Sector Weight and Turnover Constraints

László PÁL

Sapientia Hungarian University of Transylvania (Cluj-Napoca, Romania),
Department of Economic Sciences
e-mail: pallaszlo@uni.sapientia.ro

Csongor KRISTÓ

Sapientia Hungarian University of Transylvania (Cluj-Napoca, Romania),
Department of Economic Sciences
kristolcsongor@uni.sapientia.ro

Abstract. This study examines the index tracking problem using constrained optimization with sector weight and turnover constraints. We implement a mixed-integer quadratic programming (MIQP) model using the Gurobi framework. Our goal is to construct concentrated portfolios that replicate the Standard & Poor's (S&P) 100 performance. The approach employs a rolling window methodology with in-sample optimization and out-of-sample validation periods to assess performance across varying market conditions. The analysis evaluates constraint combination effects on tracking accuracy, portfolio stability, and computational requirements for 10- and 20-stock portfolios. Results demonstrate that constraints improve index alignment and portfolio stability, while enabling substantial computational efficiency gains primarily through solution space reduction.

Keywords: index tracking, portfolio optimization, mixed-integer programming, Gurobi solver

JEL Classification: G11, C61

1. Introduction

Index tracking represents a cornerstone of passive portfolio management, where the primary objective is to replicate the market benchmark's performance while minimizing costs and operational complexity. Passive investment strategies have gained substantial investor interest because the Efficient Market Hypothesis suggests that consistently outperforming the market through active management

is extraordinarily difficult. Silva et al. (2023) provide a comprehensive systematic review of solution approaches for the index tracking problem, demonstrating the diversity of methodological developments in this field.

Full replication involves holding all index securities in their exact weights, but this approach becomes impractical for broad indices. Beasley et al. (2003) demonstrated that full replication creates excessive transaction costs, administrative burdens, and liquidity constraints for most investors. Computational complexity increases further when managing hundreds of positions with continuous rebalancing requirements.

Index tracking optimization incorporates practical constraints, which significantly increases problem complexity. Canakgoz and Beasley (2008) showed that mixed-integer programming formulations are necessary to handle cardinality constraints that limit portfolio size. This combinatorial aspect makes the problem NP-hard, requiring sophisticated computational approaches.

In real-world portfolio management, additional constraints are required beyond simple cardinality limits. Sector constraints ensure diversified representation across different economic sectors, as demonstrated by Canakgoz and Beasley (2008) and Wu et al. (2017). The latter study revealed that unconstrained portfolios concentrate in only 5 sectors when tracking the S&P 500's 10 economic sectors. Turnover constraints control transaction costs by limiting portfolio changes between rebalancing periods. Beasley et al. (2003) implemented proportional transaction costs with systematic revision, while Strub and Baumann (2018) emphasized that turnover constraints are essential for managing the trade-off between transaction costs and tracking accuracy, particularly during volatile market conditions.

The computational complexity of constrained index tracking has motivated diverse solution approaches. Heuristic methods have gained prominence for large-scale problems. Genetic algorithms (García et al., 2018; Sant'Anna et al., 2017) and simulated annealing (Rubio-García et al., 2024) show competitive performance for cardinality-constrained problems. Mutunge and Haugland (2018) developed systematic algorithms specifically focused on minimizing tracking error under cardinality constraints. Sparse optimization techniques offer an alternative paradigm, where regularization methods achieve implicit asset selection while maintaining tractable linear programming formulations (Benidis et al., 2018). More recently, deep learning approaches have emerged for enhanced index tracking, with neural networks capturing complex market dynamics and regime-switching behaviour to improve tracking performance (Dai and Li, 2024). Anis and Kwon (2025) proposed an end-to-end, decision-based framework for cardinality-constrained portfolio optimization that integrates the entire optimization pipeline.

Despite these methodological advances, the interaction between sector weight deviation constraints and turnover limits within MIQP frameworks remains underexplored. Wu et al. (2017) developed a cluster-based approach combining sector cardinality constraints with turnover limits, where constraints control the number of assets per sector. Strub and Baumann (2018) examined turnover constraints in tracking error models. However, limited research has investigated the interaction between sector weight deviation constraints and turnover limits regarding their combined effects on optimization performance and computational efficiency.

This study develops a comprehensive MIQP framework integrating sector weight deviation and turnover constraints for S&P 100 tracking. Our primary objective is to demonstrate how the simultaneous implementation of sector weight and turnover constraints improves index tracking quality through enhanced sector alignment and portfolio stability, while the constraint-induced solution space reduction enables significant computational efficiency gains, particularly for concentrated portfolios. We use the Gurobi solver with rolling window methodology to systematically evaluate constraint interactions across 10- and 20-stock portfolios, analysing tracking accuracy, portfolio stability, and computational feasibility.

The paper proceeds as follows: Section 2 presents our MIQP model, Section 3 describes the methodology and data, Section 4 analyses empirical results, and Section 5 presents conclusions and outlines future research directions.

2. Model Formulation

This section develops the mathematical framework for constrained index tracking optimization. The formulation progresses from a basic cardinality-constrained model to a comprehensive MIQP framework incorporating sector weight and turnover constraints.

2.1. Basic Model

The mathematical model employs return data r_{it} representing the performance of stock i during period t , where $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\}$. The benchmark index returns are denoted by R_t for each time period. The optimization framework utilizes T historical observations across N available securities. The portfolio composition is determined by continuous weight variables, w_i representing the allocation to each asset i . Binary selection indicators z_i equal 1 when asset i is selected for inclusion and 0 otherwise. The cardinality limit K restricts the maximum number of securities in the final portfolio.

The cardinality constrained index tracking problem minimizes the sum of squared tracking errors, as formulated by García et al. (2018) and Rubio-García et al. (2024):

$$\min_{w,z} \frac{1}{T} \sum_{t=1}^T \left(\sum_{i=1}^N w_i r_{it} - R_t \right)^2 \quad (1)$$

subject to:

$$\sum_{i=1}^N w_i = 1 \quad (2)$$

$$\sum_{i=1}^N z_i = K \quad (3)$$

$$w_i \leq z_i, \quad \forall i \in \{1, \dots, N\} \quad (4)$$

$$w_i \geq 0, \quad \forall i \in \{1, \dots, N\}, \quad (5)$$

$$z_i \in \{0,1\}, \quad \forall i \in \{1, \dots, N\}, \quad (6)$$

This formulation represents an MIQP problem, which is NP-hard due to the combinatorial nature of asset selection combined with the quadratic objective function. The quadratic objective implicitly captures asset correlations through cross-product terms in the expanded form, penalizing positively correlated assets and rewarding negatively correlated ones. While this differs from explicit covariance matrix formulations, diversification benefits emerge naturally through historical return patterns, supplemented by the sector and turnover constraints described below.

2.2. Sector Constraints

To control sector exposure, we introduce sector weight constraints. Wu et al. (2017) demonstrated the necessity of sector constraints, showing that unconstrained cluster-based portfolios concentrate in only 5 sectors. While their approach constrains the number of assets selected per sector, we directly limit sector weight deviations from the benchmark to maintain tighter alignment with index sector allocations. Let I_s be the set of assets belonging to sector $s \in \{1, \dots, S\}$, t_s the target weight for sector s , and δ the maximum allowed deviation. We impose lower (7) and upper (8) bounds on sector weights:

$$\sum_{i \in I_s} w_i \geq t_s - \delta, \forall s \in \{1, \dots, S\} \quad (7)$$

$$\sum_{i \in I_s} w_i \leq t_s + \delta, \forall s \in \{1, \dots, S\} \quad (8)$$

These constraints ensure that the portfolio sector allocation remains within δ percentage points of the target sector weights t_s , preventing excessive concentration in any single sector and maintaining diversification characteristics similar to the benchmark index.

2.3. Turnover Constraints

To limit portfolio rebalancing activity, we introduce turnover constraints, as emphasized by Strub and Baumann (2018). Let w_i^{prev} denote the previous period weight of asset i . The turnover constraint with maximum limit τ is:

$$\sum_{i=1}^N |w_i - w_i^{prev}| \leq \tau \quad (9)$$

To transform this into a linear constraint, we introduce auxiliary variables u_i^+ and u_i^- representing purchases and sales respectively. The absolute value constraint can be linearized as:

$$w_i - w_i^{prev} = u_i^+ - u_i^-, \forall i \in \{1, \dots, N\} \quad (10)$$

$$\sum_{i=1}^N (u_i^+ + u_i^-) \leq \tau \quad (11)$$

$$u_i^+ \geq 0, u_i^- \geq 0, \forall i \in \{1, \dots, N\} \quad (12)$$

This formulation controls transaction costs by ensuring the sum of all absolute weight changes does not exceed τ , promoting portfolio stability and reducing trading activity. Wu et al. (2017) note that such constraints are critical for practical portfolio management, particularly in volatile markets where excessive rebalancing can erode returns.

3. Methodology and Data

This section details the solution methodology and empirical framework used to assess the performance and computational efficiency of constrained MIQP optimization for index tracking. The experimental design enables the comprehensive evaluation of constraint interactions and their impact on both solution quality and computational efficiency.

3.1. Solution Approach

The optimization problem (1)-(12) is solved using the Gurobi Optimizer,¹ a state-of-the-art solver that implements a branch-and-bound algorithm tailored for mixed-integer quadratic programming problems. The solver systematically explores the solution space by partitioning the feasible region through branching on integer variables and employs advanced cutting-plane techniques to tighten the relaxation bounds. A 10-minute time limit is imposed on each optimization instance to balance solution quality and runtime. If the solver reaches the time limit, it returns the best feasible solution found and an optimality gap bound. A mixed-integer programming gap (MIP gap) of 0% indicates that the solver has proven optimality within its numerical tolerances. When the solver reaches the time limit before achieving a 0% gap, it returns the best feasible solution found along with the current optimality gap, indicating that the solution may be suboptimal. This approach allows for efficient approximation in complex, large-scale problem instances where exact solutions may be computationally prohibitive.

The computational environment consists of an Intel Xeon Gold 5415 processor operating at 2.9 GHz with 128 GB RAM. The implementation uses Python with Gurobi 12 for the optimization framework.

3.2. Data and Experimental Setup

The empirical study uses S&P 100 index data spanning from 2019 to 2024 with daily frequency, comprising a total of 101 stocks. The S&P 100 represents a computationally manageable benchmark that provides adequate complexity for testing sector and turnover constraints effectively, while offering the superior liquidity and market depth characteristic of US equity markets. This period captures diverse market conditions and significant changes in sector composition over time. Market capitalization data for S&P 100 constituents was downloaded from Financial Modeling Prep.² We calculated sector weights using a market

¹ Available at: <https://www.gurobi.com>.

² <https://www.financialmodelingprep.com>.

capitalization-weighted approach, where each sector's weight represents its percentage of total index market capitalization. This process was repeated for each trading day to create time series of sector weights spanning the period of 2019–2024, ensuring that sector constraints in the optimization model reflect actual market-based sector allocations.

As shown in *Table 1*, the Technology sector emerged as dominant, growing from 22.13% to 34.93%, while other major sectors showed varying patterns: Communication Services remained stable (~19%), Financial Services declined significantly (16.38% to 11.04%), and Healthcare fluctuated considerably.

Table 1. *Changes in sector weights over time*

Sector	2019	2020	2021	2022	2023	2024
Technology	22.13%	25.81%	28.28%	25.91%	32.08%	34.93%
Comm. Services	19.00%	19.21%	19.70%	15.09%	17.84%	19.03%
Financial Services	16.38%	13.21%	12.34%	13.89%	11.78%	11.04%
Consumer Cyclical	9.87%	14.13%	13.70%	9.97%	11.55%	12.28%
Healthcare	12.08%	10.76%	10.53%	14.28%	10.66%	8.46%
Consumer Defensive	8.69%	7.64%	6.77%	8.41%	6.66%	6.45%
Industrials	5.91%	5.18%	4.43%	5.69%	4.62%	3.88%
Energy	3.07%	1.61%	1.87%	3.98%	2.65%	2.17%
Utilities	1.31%	1.19%	1.08%	1.37%	0.90%	0.82%
Basic Materials	0.80%	0.74%	0.70%	0.84%	0.77%	0.58%
Real Estate	0.76%	0.53%	0.60%	0.58%	0.48%	0.36%

The methodology employs a rolling window approach with an in-sample period of 250 days for optimization and an out-of-sample period of 120 days for performance evaluation, resulting in 10 optimization windows throughout the study period. Portfolio weights from the previous optimization serve as initial values (warm start) for the current optimization.

The study examines concentrated portfolios with 10 stocks (10% of S&P 100 constituents) and moderately diversified portfolios with 20 stocks (20% of constituents). These sizes align with index tracking literature, where small portfolios typically represent 2–20% of benchmark constituents (Wu et al., 2017; Mutunge and Haugland, 2018), allowing systematic evaluation of constraint effectiveness across different concentration levels.

The experimental framework combines sector and turnover constraints to create a comprehensive constraint environment. Sector constraints employ deviation parameters $\delta = 0.02, 0.03, 0.04$, representing 2%, 3%, and 4% maximum deviations from benchmark sector weights, which provide reasonable flexibility while maintaining sector discipline. Turnover constraints use limits $\tau = 0.2, 0.3$,

0.4, corresponding to 20%, 30%, and 40% maximum portfolio turnover between rebalancing periods, ranging from typical index fund levels (20%) to more active management approaches (40%). Combined with the baseline no-constraints case, this design yields 10 different constraint scenarios for systematic evaluation.

Performance evaluation employs a set of metrics to assess portfolio effectiveness across multiple dimensions. *Net return* measures the annualized portfolio return after accounting for transaction costs, while *cumulative return* captures the total return over the entire evaluation period. *Volatility* quantifies the standard deviation of portfolio returns, providing insight into risk levels. *Tracking error* (TE) measures the standard deviation of excess returns between the portfolio and benchmark, indicating replication precision. The *Sharpe ratio* evaluates risk-adjusted performance relative to the risk-free rate. *Maximum drawdown* (MDD) assesses the largest peak-to-trough decline, capturing downside risk exposure. *Average sector deviation* measures the mean absolute deviation from benchmark sector weights, indicating sector allocation precision. *Average turnover* quantifies the mean portfolio rebalancing activity across all periods. Transaction costs are set at 0.2% per trade, reflecting typical institutional brokerage fees for equity transactions. The risk-free rate is set at 2.5% annually, representing the average rate observed during the 2019–2024 study period.

4. Results and Analysis

This section analyses portfolio performance and computational efficiency results obtained from the rolling window methodology across all constraint combinations. The evaluation covers both 10-stock and 20-stock portfolios, assessing the effects of different constraint combinations on index tracking effectiveness and computational requirements across the 10 rolling windows.

4.1. Results Considering 10 Stocks

Portfolio Performance

The analysis of constrained optimization results reveals how sector and turnover constraints enhance portfolio characteristics through improved sector alignment and reduced trading activity. The unconstrained baseline achieves 18.75% net return with 115.25% cumulative return, but exhibits significant deviation from index characteristics that constrained strategies systematically address. Several constrained variants achieve superior index alignment compared to the unconstrained approach. The $\delta = 0.02$, $\tau = 0.4$ strategy achieves 22.24%

volatility, which is much closer to the S&P 100's 21.97% than the unconstrained portfolio's 22.96%. Most notably, the $\delta = 0.03$, $\tau = 0.2$ combination delivers a 4.83% tracking error that improves upon the unconstrained baseline's 4.85%, while its 0.64 Sharpe ratio closely approaches the index's 0.61 level. As shown in *Table 2*, constrained strategies maintain competitive tracking errors while delivering improvements in sector discipline and portfolio stability.

Table 2. *Performance of 10-stock portfolios with combined constraints*

Strategy	Net Return (%)	Cum. Return (%)	Volatil- ity (%)	Sharpe Ratio	MDD (%)	TE (%)	Avg. Sect. Dev. (%)	Avg. Turn (%)
No Constraints	18.75	115.25	22.96	0.71	-34.23	4.85	14.22	100.79
$\delta = 0.02$, $\tau = 0.2$	16.54	94.92	22.43	0.63	-32.03	5.39	2.00	20.00
$\delta = 0.02$, $\tau = 0.3$	13.46	68.18	22.50	0.49	-32.03	5.38	2.00	28.82
$\delta = 0.02$, $\tau = 0.4$	16.90	98.67	22.24	0.65	-32.03	5.10	2.00	38.73
$\delta = 0.03$, $\tau = 0.2$	17.01	98.49	22.83	0.64	-33.03	4.83	3.00	20.00
$\delta = 0.03$, $\tau = 0.3$	14.50	76.55	22.61	0.53	-33.03	5.04	3.00	29.68
$\delta = 0.03$, $\tau = 0.4$	16.45	94.20	22.38	0.62	-33.03	4.86	3.00	39.63
$\delta = 0.04$, $\tau = 0.2$	16.29	91.76	22.84	0.60	-33.26	4.86	3.94	20.00
$\delta = 0.04$, $\tau = 0.3$	15.97	88.85	22.82	0.59	-33.26	4.91	3.99	29.85
$\delta = 0.04$, $\tau = 0.4$	15.25	83.32	22.41	0.57	-33.26	4.66	3.96	39.50
S&P 100 Index	15.97	90.56	21.97	0.61	-31.53	N/A	N/A	N/A

The tracking error results align with established findings for concentrated portfolios. Scozzari et al. (2013) achieved 5.75% tracking error for 10-stock Nikkei 225 portfolios, confirming that concentrated tracking typically produces tracking errors in the 4–6% range. While Mutunge and Haugland (2018) reported 2.77% for 10-stock S&P 500 portfolios, their covariance-based ex-ante definition differs fundamentally from our ex-post realized return standard deviation approach.

The cumulative return comparison in *Figure 1* demonstrates the evolution of tracking quality over the study period. Until early 2022, all cumulative returns closely follow the index trajectory, indicating effective tracking across different constraint combinations. After 2022, deviations become more pronounced, with the unconstrained case showing the largest divergence from the benchmark. At the final evaluation point, some constrained strategies perform slightly below the index curve, while others achieve marginally better results, with the $\delta = 0.04$, $\tau = 0.2$ strategy showing a particularly favourable final positioning relative to the benchmark.




Figure 1. Cumulative return comparison for 10-stock portfolios

Sector constraints demonstrate effectiveness in controlling allocation drift, reducing average sector deviation from the unconstrained portfolio's excessive 14.22% to a disciplined 2–4% range across all constrained variants. This improvement directly enhances index tracking quality by preventing excessive sector concentrations that compromise benchmark replication precision. The $\delta = 0.03, \tau = 0.2$ strategy exemplifies this balance, achieving tight sector control (3.00% deviation) while maintaining competitive 17.01% net returns.

The portfolio composition analysis in *Figure 2* reveals the stabilizing effect of constraints on stock selection patterns. This figure displays the evolution of the top 10 holdings as stacked area plots, where the top holdings are determined by calculating the average weight of each stock across all rebalancing dates and selecting the 10 stocks with the highest average weights. The unconstrained portfolio exhibits volatile allocation patterns with frequent changes in major holdings, while constrained variants maintain notably more stable stock compositions over time. Particularly stable holdings across most strategies include major technology stocks, such as Apple (AAPL) and Microsoft (MSFT), which consistently maintain substantial positions due to their prominent index weights and strong performance characteristics.

Turnover constraints address critical implementation challenges by reducing the unconstrained portfolio's excessive 100.79% turnover to manageable 20–40% levels. This reduction improves net returns through lower transaction costs while enhancing portfolio stability. Strub and Baumann (2018) demonstrated that managing the trade-off between transaction costs and tracking similarity is fundamental to index fund optimization. Strategies with turnover limits of $\tau = 0.2$

consistently produce the most stable tracking performance across different sector constraint levels.




Figure 2. Portfolio composition changes for 10-stock portfolios

Computational Efficiency

The logarithmic scale plot in *Figure 3* demonstrates the computational challenge of Window 1, where all constraint combinations require between 50 and 600 seconds, with the $\delta = 0.04, \tau = 0.4$ strategy experiencing the most severe computational burden at 600 seconds. Window 1 represents a universal computational bottleneck across all strategies, as the solver lacks prior solution information to initialize the optimization. From Window 2 onwards, all constrained strategies achieve remarkable computational convergence, dropping to consistently under 1 second solve times across all constraint combinations. This reduction spans 2-3 orders of magnitude, primarily attributable to the turnover constraints that substantially reduce the solution space by limiting portfolio composition

changes between consecutive periods. The MIP gap results confirm this computational pattern, with Window 1 achieving only partial optimal solutions, while Windows 2–10 consistently achieve optimal solutions with zero gaps across all constraint combinations. This computational efficiency transformation enables practical real-time portfolio management, as subsequent optimizations require minimal computational resources.




Figure 3. Computational time evolution for 10-stock portfolio optimization

4.2. Results Considering 20 Stocks

Portfolio Performance

The analysis of 20-stock constrained optimization results reveals enhanced index tracking capabilities and substantially different constraint dynamics compared to the 10-stock case (see *Table 3*). The unconstrained baseline achieves 20.34% net return with 135.39% cumulative return, demonstrating superior absolute performance while maintaining significantly better index alignment characteristics than its 10-stock counterpart.

The 20-stock portfolios show improved index tracking quality across all strategies. The unconstrained portfolio achieves 3.13% tracking error, representing a substantial

35% improvement over the 10-stock case's 4.85% while maintaining 21.69% volatility very close to the S&P 100's 21.97%. This improvement demonstrates that doubling the portfolio size from 10 to 20 stocks provides substantial diversification benefits, enabling closer index replication through the better representation of benchmark characteristics.

Table 3. *Performance of 20-stock portfolios with combined constraints*

Strategy	Net Return (%)	Cum. Return (%)	Volatil- ity (%)	Sharpe Ratio	MDD (%)	TE (%)	Avg. Sect. Dev. (%)	Avg. Turn (%)
No Constraints	20.34	135.39	21.69	0.82	-30.12	3.13	7.58	92.07
$\delta = 0.02, \tau = 0.2$	17.08	101.02	21.95	0.66	-32.27	3.19	2.00	20.00
$\delta = 0.02, \tau = 0.3$	17.62	105.87	22.12	0.68	-32.27	3.14	2.00	30.00
$\delta = 0.02, \tau = 0.4$	17.70	106.81	22.07	0.69	-32.27	3.13	2.00	40.00
$\delta = 0.03, \tau = 0.2$	16.91	99.46	21.91	0.66	-30.23	3.40	2.92	20.00
$\delta = 0.03, \tau = 0.3$	18.71	117.13	21.96	0.74	-30.23	3.39	3.00	30.00
$\delta = 0.03, \tau = 0.4$	18.25	112.51	21.95	0.72	-30.23	3.30	3.00	40.00
$\delta = 0.04, \tau = 0.2$	16.05	91.50	21.90	0.62	-32.34	3.26	4.00	20.00
$\delta = 0.04, \tau = 0.3$	17.55	105.46	22.01	0.68	-32.34	3.21	3.82	30.00
$\delta = 0.04, \tau = 0.4$	16.51	95.25	22.11	0.63	-32.34	3.09	3.78	39.78
S&P 100 Index	15.97	90.56	21.97	0.61	-31.53	N/A	N/A	N/A

Constrained strategies maintain competitive tracking performance while delivering enhanced implementation characteristics. Several constrained strategies achieve tracking errors of 3.09–3.40%, closely matching the unconstrained baseline's 3.13%. Notably, the $\delta = 0.02, \tau = 0.4$ combination delivers 3.13% tracking error, exactly matching the unconstrained performance while maintaining disciplined sector control (2.00% sector deviation vs 7.58% unconstrained). These constrained tracking error results align well with established findings for portfolios of comparable size. Scozzari et al. (2013) reported 3.41% tracking error for constrained 20-stock Nikkei 225 portfolios under UCITS constraints, confirming that 20-stock constrained portfolios consistently achieve tracking errors in the 3–4% range.

The 20-stock portfolios also demonstrate superior risk-adjusted performance. The unconstrained portfolio achieves a Sharpe ratio of 0.82, exceeding the index's 0.61 level by 35%, while several constrained strategies achieve Sharpe ratios of 0.62–0.74, maintaining competitive risk-adjusted performance while delivering the practical benefits of reduced turnover and controlled sector exposure. As shown in Table 3, 20-stock portfolios consistently outperform 10-stock equivalents across multiple dimensions, including tracking error, volatility alignment, sector discipline, and portfolio stability.

The cumulative return comparison presented in *Figure 4* demonstrates the superior tracking characteristics of the 20-stock portfolios over the study period. The cumulative return trajectories show that all strategies position slightly above the index but maintain remarkably tight clustering without large dispersions. Unlike the 10-stock case where significant deviations emerged after 2022, the 20-stock portfolios maintain much closer proximity to the benchmark throughout the entire period. At the final evaluation point, the differences between strategies remain moderate. The $\delta = 0.04, \tau = 0.2$ strategy delivers cumulative returns closest to the index, making it the optimal choice for pure index tracking.




Figure 4. Cumulative return comparison for 20-stock portfolios

Contrasting with 10-stock results, sector constraints in 20-stock portfolios impose significantly lower performance penalties. The unconstrained portfolio's 7.58% sector deviation is much more manageable than the 10-stock case's 14.22%, and constrained strategies achieve 2–4% deviations with substantially smaller return sacrifices. The $\delta = 0.03, \tau = 0.3$ strategy achieves 18.71% net returns, which is substantially higher than any 10-stock constrained strategy. At the same time, it maintains excellent sector control at 3.00%. This demonstrates that increased diversification naturally improves sector balance and reduces the trade-off between constraint compliance and performance. Wu et al. (2017) similarly demonstrated that sector limit constraints improve diversification, with their sector-constrained portfolios achieving more balanced sector allocations and better tracking stability compared to unconstrained alternatives.

The portfolio composition analysis in *Figure 5* shows similar patterns to the 10-stock case, with constrained variants maintaining stable exposures that mirror index composition. However, the weight changes are even smoother than in the 10-stock case, with major holdings like Apple (AAPL) and Microsoft (MSFT) displaying exceptionally stable positions throughout the evaluation period.




Figure 5. Portfolio composition changes for 20-stock portfolios

Computational Efficiency

The computational landscape for 20-stock portfolios presents different challenges compared to the 10-stock case, as illustrated in *Figure 6*. This escalation in computational requirements aligns with the well-documented NP-hard complexity of cardinality-constrained index tracking (Canakgoz and Beasley, 2008; Mutunge and Haugland, 2018), where solution times grow substantially as portfolio size increases. Despite increased computational complexity for 20-stock portfolios, the turnover constraints continue to provide meaningful solution space reduction. While computational times are indeed larger, ranging from 10 to 600 seconds across Windows 2–10 compared to the performance of the 10-stock case, the optimization method demonstrates robust convergence in most cases.




Figure 6. Computational time evolution for 20-stock portfolio optimization

Table 4 presents detailed MIP gap analysis for the 20-stock optimization across the 10 rolling windows. For each window, we report statistics across all 9 constraint strategies: the count of optimal solutions (zero MIP gap), the number of instances hitting the 600-second time limit, and the mean, standard deviation, minimum, and maximum MIP gaps for non-optimal solutions.

Table 4. MIP gap results for 20-stock portfolio optimization

Window	Optimal Count	Time Limit	Mean Gap (%)	Std. Gap (%)	Min. Gap (%)	Max. Gap (%)
1	0	9	78.90	2.90	75.80	82.80
2	9	0	0.00	0.00	0.00	0.00
3	7	2	1.50	3.00	0.00	8.50
4	6	3	2.10	4.30	0.00	11.00
5	8	1	1.50	4.40	0.00	13.20
6	8	1	1.00	2.90	0.00	8.60
7	8	1	0.00	0.10	0.00	0.30
8	7	2	2.90	5.90	0.00	15.80
9	8	1	0.00	0.00	0.00	0.00
10	9	0	0.00	0.00	0.00	0.00

Window 1 computational burden intensifies significantly, with all constraint combinations requiring 600 seconds and hitting the time limit, as confirmed by *Table 4* showing zero optimal solutions achieved and MIP gaps averaging 78.9%. This represents a substantial deterioration from the 10-stock case, where at least partial optimal solutions were obtained in Window 1.

Despite the increased computational times, the method shows strong convergence patterns: Windows 2, 9, and 10 achieve all 9 optimal solutions, while Windows 3–8 obtain 6–8 optimal solutions out of 9 constraint combinations. Even in non-convergent cases, the MIP gap values remain small (typically under 3%, maximum 15.80%), indicating that near-optimal solutions are consistently achieved within the time limit. *Figure 6* reveals that tighter constraints ($\delta = 0.02$ combinations) generally achieve more stable computational performance, while the $\delta = 0.03, \tau = 0.2$ strategy demonstrates particularly consistent efficiency after Window 2. Notably, strategies with lower turnover limits ($\tau = 0.2$) consistently exhibit faster computational times than their higher turnover counterparts ($\tau = 0.3, \tau = 0.4$), directly reflecting the reduced solution space imposed by stricter turnover constraints. The persistent computational volatility in some strategies reflects the increased complexity of the 20-stock optimization landscape rather than fundamental algorithmic failure. The turnover constraints continue to reduce the solution space after Window 1, though less dramatically than in the 10-stock case. The combination of reasonable solve times and consistently small MIP gaps demonstrates that the method remains practically viable for 20-stock portfolio optimization, delivering high-quality solutions within acceptable computational bounds.

5. Conclusions

This study demonstrates the effectiveness of constrained optimization approaches for index tracking portfolio construction using mixed-integer quadratic programming enhanced with sector weight and turnover constraints. The Gurobi optimization framework proves capable of handling complex constraint combinations efficiently.

The empirical evidence supports constrained optimization strategies. Constrained variants improve index alignment, reducing sector deviation from 14.22% to 2–4% for 10-stock portfolios and from 7.58% to 2–4% for 20-stock portfolios. Importantly, several constrained strategies achieve comparable or lower tracking errors than unconstrained approaches (e.g. 10-stock $\delta = 0.03, \tau = 0.2$: 4.83% vs 4.85% unconstrained). Turnover constraints reduce excessive portfolio turnover from ~100% to manageable 20–40% levels, improving net returns through lower transaction costs.

Constraint-induced solution space reduction yields different computational profiles by portfolio size. For 10-stock portfolios, the computational burden reduces from

hundreds of seconds to consistently under one second after the initial window, representing a 2-3 order of magnitude improvement. For 20-stock portfolios, solve times range from 10 to 600 seconds, but the method maintains practical viability with MIP gaps typically under 3% and superior performance for $\tau = 0.2$ strategies.

Our empirical analysis indicates that concentrated portfolios with relatively few assets can effectively replicate market index performance when properly optimized. Both 10-stock and 20-stock portfolios achieve competitive tracking errors, with 20-stock portfolios reaching 3.13% while maintaining superior constraint compliance. The rolling window methodology confirms consistent tracking quality across diverse market conditions spanning 2019–2024, with constrained strategies positioning close to the benchmark throughout the evaluation period.

The demonstrated effectiveness with the S&P 100 suggests promising applicability to broader indices such as the S&P 500. Future research should explore alternative techniques to further reduce solution times. The Gurobi solution pool functionality presents compelling opportunities for multi-criteria portfolio selection beyond simple tracking error minimization. Enhanced index tracking strategies that target controlled outperformance while maintaining low tracking error represent a natural extension of this constraint framework.

References

- Anis, H. T.; Kwon, R. H. (2025). End-to-end, decision-based, cardinality-constrained portfolio optimization. *European Journal of Operational Research* 320: 739–753.
- Beasley, J. E.; Meade, N.; Chang, T.-J. (2003). An evolutionary heuristic for the index tracking problem. *European Journal of Operational Research* 148: 621–643.
- Benidis, K.; Feng, Y.; Palomar, D. P. (2018). Sparse portfolios for high-dimensional financial index tracking. *IEEE Transactions on Signal Processing* 66: 155–170.
- Canakgoz, N. A.; Beasley, J. E. (2008). Mixed-integer programming approaches for index tracking and enhanced indexation. *European Journal of Operational Research* 196: 384–399.
- Dai, Z.; Li, L. (2024). Deep learning for enhanced index tracking. *Quantitative Finance* 24: 569–591.
- García, F.; Guijarro, F.; Oliver, J. (2018). Index tracking optimization with cardinality constraint: A performance comparison of genetic algorithms and tabu search heuristics. *Neural Computing and Applications* 30(8): 2625–2641.
- Mutunge, P.; Haugland, D. (2018). Minimizing the tracking error of cardinality constrained portfolios. *Computers & Operations Research* 90: 33–41.
- Rubio-García, Á.; Fernández-Lorenzo, S.; García-Ripoll, J. J.; Porras, D. (2024). Accurate solution of the index tracking problem with a hybrid simulated

- annealing algorithm. *Physica A: Statistical Mechanics and Its Applications* 639: 129637.
- Sant'Anna, L. R.; Filomena, T. P.; Guedes, P. C.; Borenstein, D. (2017). Index tracking with controlled number of assets using a hybrid heuristic combining genetic algorithm and non-linear programming. *Annals of Operations Research* 258: 849–867.
- Scozzari, A.; Paterlini, S.; di Zio, S. (2013). Exact and heuristic approaches for the index tracking problem with UCITS constraints. *Annals of Operations Research* 205: 235–250.
- Silva, J. C. S.; Silva, D. F. L.; Ferreira, L.; Almeida-Filho, A. T. (2023). A systematic literature review on solution approaches for the index tracking problem. *Annals of Operations Research* 331: 1001–1042.
- Strub, O.; Baumann, P. (2018). Optimal construction and rebalancing of index-tracking portfolios. *European Journal of Operational Research* 264: 370–387.
- Wu, D.; Kwon, R. H.; Costa, G. (2017). A constrained cluster-based approach for tracking the S&P 500 index. *International Journal of Production Economics* 193: 222–243.