International Journal of
Financial Studies

MDPI

Article

A Hybrid Genetic Algorithm with Learning-to-Rank-to-Optimization for US Equity Portfolio Construction

Ferdinantos Kottas [1,2,3]

[1] School of Business, National University of Ireland Maynooth, W23 F2H6 Maynooth, Ireland; ferdinantos.kottas.2019@mumail.ie
[2] School of Engineering, Civil Engineering, Democritus University of Thrace, 67100 Xanthi, Greece
[3] Department of Computer Science, Aristotle University of Thessaloniki, 54124 Thessaloniki, Greece

Abstract

This study develops and evaluates an equity selection pipeline that converts quarterly fundamentals into a monthly frequency, constructs profitability, leverage, liquidity, and growth characteristics, and learns a linear ranking model via a genetic algorithm (GA). The GA is trained to maximize either (i) mean monthly NDCG@30 using 12-tile relevance labels or (ii) mean monthly Spearman information coefficient (IC). The learned ranker is tested out-of-sample using monthly forward returns, benchmarked against the S&P 500, with different types of allocation weights, and further evaluated under sector concentration limits. In the last layer, the monthly-selected stock universe is used in a daily dynamic allocation which is solved by the penalized Max-Sharpe or Min-Variance optimization problems under only long positions and transaction fees. Performance is examined across Pre-COVID, COVID, Post-COVID (Train), and Final Test regimes, demonstrating how ranking objectives and diversification constraints impact performance and stability. Results show that TTM-based accounting signals, when optimized through genetic learning and disciplined allocation, yield economically meaningful stock selection and robust portfolio performance across market regimes.

Keywords: portfolio optimization; genetic algorithms; learning to rank; Sharpe ratio; mean–variance; portfolio allocation

Check for updates

Academic Editor: Indranil SenGupta

Received: 2 March 2026
Revised: 23 March 2026
Accepted: 31 March 2026
Published: 4 April 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Equity markets generate vast amounts of accounting information, yet translating this information into robust portfolio performance remains an unresolved challenge. While decades of empirical asset pricing research demonstrate that firm characteristics, such as profitability, investment intensity, leverage, liquidity, and growth, predict cross-sectional differences in expected returns, practical implementation often fails to translate statistical predictability into actionable insights. Additionally, traditional approaches often treat return prediction as a simple regression or classification task; however this overlooks the critical insight that portfolio construction benefits from learning-to-rank frameworks, which explicitly optimize for the relative ordering of assets rather than absolute return predictions. By focusing on ranking, investors align model objectives with actual investment decisions: selecting top-performing stocks for inclusion rather than minimizing prediction errors across all stocks (Kwiatkowski & Chudziak, 2025). Contemporary asset-ranking techniques using regression or classification heuristics are demonstrably suboptimal because they

fail to capture the pairwise and listwise relationships across instruments (Poh et al., 2021). Modern machine learning ranking algorithms, by contrast, learn these broader cross-asset linkages, producing substantially more accurate rankings and improving portfolio performance. Our research has demonstrated a strategy design that underlies these issues.

Financial statements are released quarterly, and their informational content becomes available only upon public filing. Many empirical studies implicitly align accounting variables with fiscal period-end dates, which introduces look-ahead bias and artificially increases the signal refresh frequency. A strategy that uses information before it is publicly available overstates its true predictive power. Realistic implementation requires strict alignment of accounting variables to official filing dates and careful construction of trailing-twelve-month (TTM) measures to smooth seasonality without violating information timing.

Furthermore, even when timing is handled correctly, the statistical objective used to convert accounting variables into forecasts is often misaligned with portfolio construction. As mentioned, the most common approach is to treat stock selection as a regression or classification task, minimizing squared errors or cross-entropy loss across the full cross-section. Machine learning methods have been shown to enhance stock price predictions compared to traditional financial models by capturing nonlinear relationships between factors and returns (Agarwalla, 2024), making them particularly valuable for extracting value from fundamental accounting signals. However, investors do not allocate capital proportionally to predicted numerical returns. Instead, they select a subset of securities expected to outperform. The economic objective is, therefore, to rank securities correctly—especially those at the top of the distribution that will receive capital allocation. A model that minimizes prediction error may still misorder the highest-return stocks, leading to weak portfolio performance.

Recent research in machine learning emphasizes that learning-to-rank frameworks better align with the economic structure of portfolio selection. Rather than predicting absolute return magnitudes, ranking models directly optimize the relative ordering of securities. Empirical evidence suggests that ranking-based systematic strategies can materially outperform regression-based approaches, in some cases, producing significantly higher Sharpe ratios and more stable performance (Poh et al., 2021). Furthermore, modern architectures that directly optimize ranking metrics, such as Normalized Discounted Cumulative Gain (NDCG), demonstrate that aligning the training objective with portfolio inclusion decisions improves risk-adjusted returns (Sakagawa & Mori, 2022).

However, even if ranking is performed correctly, portfolio implementation introduces additional challenges. High-scoring stocks may cluster within a small number of sectors, embedding implicit industry bets that dominate firm-level fundamentals. Moreover, portfolio weights are sensitive to estimation error, particularly in short-horizon mean forecasts. The mean-variance literature shows that estimation error can substantially degrade out-of-sample performance if not properly regularized (DeMiguel et al., 2009). Robust portfolio construction must, therefore, address both diversification and allocation stability.

These issues became particularly evident during the COVID-19 pandemic. The crisis introduced extreme volatility, sectoral dispersion, and structural breaks in economic activity (Baker et al., 2020). Strategies calibrated under stable market conditions were challenged by regime shifts and abrupt changes in cross-sectional dynamics. Any credible equity selection framework must, therefore, demonstrate robustness across heterogeneous market environments, including crisis periods and strictly out-of-sample windows.

Against this backdrop, this study develops a two-stage equity selection framework designed to address the alignment problem between accounting information, ranking objectives, and portfolio construction. Quarterly accounting fundamentals are transformed from a quarterly to a monthly frequency, aligned with filing dates, and mapped into cross-sectional monthly ranks to ensure scale comparability. A linear scoring function aggregates these

characteristics using weights optimized by a genetic algorithm (GA). Genetic algorithms are particularly suited to ranking objectives because they do not require differentiable loss functions and can directly optimize non-convex, order-dependent metrics.

Two economically distinct ranking objectives are considered. The first maximizes NDCG@30, a top-heavy metric that emphasizes correct ordering of extreme outperformers likely to enter a Top-K portfolio. The second maximizes the monthly Spearman Information Coefficient (IC), which measures monotonic alignment between predicted scores and realized forward returns across the entire cross-section. By optimizing identical accounting features under these distinct objectives, the study isolates how objective design influences concentration, stability, and economic performance.

Portfolio construction proceeds in two layers. In the first stage, the selection layer ranks stocks monthly, and Top-K portfolios are formed either without restriction or under a sector-capped rule that limits the number of selected stocks per sector. This allows evaluation of whether alpha arises from firm-level fundamentals or from concentrated industry exposures. In the second stage, the selected universe may be dynamically reweighted using daily penalized minimum-variance or maximum-Sharpe optimization under long-only and turnover constraints. This separation of the selection and allocation layers enables the decomposition of performance into cross-sectional alpha extraction and conditional risk management.

Performance is evaluated across four economically distinct regimes: Pre-COVID expansion, COVID stress, Post-COVID recovery (training window), and a strictly held-out Final Test period. By training exclusively on the post-COVID window and holding model parameters fixed thereafter, the framework enforces genuine out-of-sample validation in a non-stationary environment.

The central motivation of this study is, therefore, to bridge the gap between accounting-based asset pricing evidence and implementable portfolio construction. By combining strict information timing, ranking-consistent optimization, evolutionary search, and disciplined allocation, the paper demonstrates how fundamental signals can be transformed into a robust and economically interpretable long-only equity strategy. While the S&P 500 index is used as the primary benchmark, the study further evaluates robustness through alternative portfolio constructions, including different portfolio sizes (Top-10 to Top-50) and weighting schemes (equal and exponential). These variations provide a meaningful comparison to standard fundamental screening approaches and allow assessment of whether performance is driven by the ranking signal rather than portfolio construction choices.

2. Related Literature and Research Background

This study builds on several strands of the literature spanning empirical asset pricing, learning-to-rank methodologies, evolutionary optimization, and portfolio construction under estimation risk.

A substantial body of research documents that accounting characteristics explain systematic variation in expected returns. Fama and French (1992, 2015) show that size, value, profitability, and investment factors capture return patterns unexplained by market beta. Novy-Marx (2013) demonstrates that gross profitability scaled by assets predicts returns with magnitudes comparable to value effects. Cooper et al. (2008) show that asset growth is negatively related to subsequent returns, while Hou et al. (2015) develop a q-theory framework linking investment and profitability to expected returns. Green et al. (2017) document a large set of accounting-based characteristics providing independent cross-sectional information.

A critical methodological issue in accounting-based prediction concerns the timing of information. Bernard and Thomas (1990) show that improper modeling of earnings

announcement timing can overstate abnormal returns. Standard practice, therefore, lags accounting variables and constructs trailing-twelve-month aggregates to smooth seasonal fluctuations and better approximate the information set available to investors (Penman, 2013). Realistic timing assumptions are central to credible backtesting.

Beyond characteristic construction, the statistical objective of prediction has evolved. Traditional cross-sectional regressions minimize squared errors, implicitly treating all prediction mistakes equally. However, active equity management depends primarily on relative ranking accuracy. The Spearman Information Coefficient (IC), defined as the cross-sectional rank correlation between predicted scores and realized returns, has become a standard measure of signal quality (Grinold & Kahn, 2000). The IC captures monotonic alignment but does not directly emphasize the top of the ranking.

Learning-to-rank (LTR) methods offer a more direct alignment between optimization and portfolio construction. Originating in information retrieval (Liu, 2009), LTR frameworks optimize ordering rather than the magnitude of predictions. Listwise ranking objectives, such as NDCG place greater weight on the correct ordering of top-ranked items (Järvelin & Kekäläinen, 2002). In financial contexts, ranking-based strategies have demonstrated improved performance relative to regression-based models (Poh et al., 2021). By emphasizing economically relevant positions, ranking objectives better reflects capital allocation decisions.

Optimizing ranking metrics introduces computational challenges because such objectives depend on order statistics and are typically non-differentiable. Genetic algorithms (Goldberg, 1989) provide a flexible metaheuristic solution. Population-based evolutionary search can navigate non-convex parameter spaces without relying on gradient information. In finance, GA-based portfolio optimization has shown effectiveness under cardinality constraints and in the presence of uncertainty (Fransisca et al., 2024; Gunjan & Bhattacharyya, 2024). Hybrid evolutionary-learning frameworks further demonstrate strong cross-sectional predictive capability (Flores-Fernandez et al., 2024; Nufus et al., 2025).

Additionally, the portfolio allocation literature highlights the impact of estimation error. Markowitz (1952) establishes the mean-variance framework, but subsequent research shows that estimation error in expected returns can severely degrade out-of-sample performance (DeMiguel et al., 2009). Regularization and shrinkage methods are, therefore, widely used to stabilize optimization. Industry structure also plays a central role in return co-movement, as documented by Moskowitz and Grinblatt (1999), motivating sector-level diversification constraints.

Finally, the present study integrates these strands by combining realistic accounting timing, ranking-based optimization, evolutionary search, and constrained portfolio allocation within a unified empirical framework. By directly comparing top-focused (NDCG) and globally monotonic (IC) ranking objectives under identical accounting features, this study extends the existing literature on both asset pricing and machine learning applications in finance while evaluating robustness across structurally distinct market regimes.

3. Data Sources and Collection

This study combines firm-level accounting fundamentals with daily and monthly market data to construct a large-scale cross-sectional panel suitable for ranking-based equity selection and portfolio optimization. Particular emphasis is placed on chronological consistency and realistic information timing to ensure that all signals reflect information available to investors at the time of portfolio formation.

The full sample spans from 1 January 2017 through 31 October 2025. The choice of starting date reflects the availability of consistently formatted accounting data across firms and ensures sufficient pre-crisis observations prior to the COVID-19 shock.

Firm fundamentals are obtained from End-of-Day Historical Data (EODHD), which provides standardized quarterly financial statements for U.S.-listed companies. The dataset includes income statement, balance sheet, and cash flow statement items (see Appendix A). Each accounting observation is aligned to its official reporting (filing) date rather than the fiscal or submitted period (quarterly). A variable becomes observable only for calendar months, specifically, the filing date to the corresponding month-end. The value is then carried forward until the next reported filing. This procedure prevents look-ahead bias and reflects the fact that accounting information updates discretely rather than continuously. This procedure ensures that accounting information is not used before it is publicly available and reflects the discrete update frequency of financial statements. Because quarterly results are not continuously revised, the carry-forward mechanism realistically captures how investors observe fundamentals between reporting dates. A key transformation is the construction of trailing-twelve-month (TTM) variables for flow quantities. Let $F_{i,q}$ denote a quarterly flow variable (e.g., revenue or net income) for company i. The TTM version at quarter q is defined as:

$$F_{i,q}^{TTM} = \sum_{t=0}^{3} F_{i,q-k} \quad (1)$$

This aggregation smooths seasonal variation and approximates annual performance while preserving timeliness. TTM construction is applied prior to monthly carry-forward alignment. Only the TTM value for the most recently reported quarter is observable each month.

The final accounting panel is converted to a monthly frequency beginning January 2017. The dataset includes 143 accounting variables per firm-month observation. These variables consist of directly downloaded statement items, TTM aggregates, and constructed ratios.

Table 1 presents the total number of months and the number of unique stocks across the regimes used in our analysis. The split into several regimes serves two purposes. The first one is to separate the stressful period, in our case, the COVID period, which represents a structural break characterized by extreme volatility and macroeconomic disruption. Separating this regime allows examination of robustness under stress. Second, model training is restricted to the post-COVID period (July 2021–October 2024) when learning optimal ranking weights, ensuring that the final test period remains strictly out-of-sample. Research on market regime adaptation shows that models trained on period-by-period rolling windows can better capture time-varying relationships between fundamentals and returns compared than static training (Capponi et al., 2025). The COVID period (February 2020–June 2021) is treated as a stress regime, characterized by extreme volatility, sectoral dispersion, and rapid changes in economic conditions. This period provides a natural stress test for evaluating the robustness of the proposed framework under non-stationary market dynamics.

Table 1. Total number of months, and unique US stocks across the regimes.

Regime	Period	Months	Unique Firms
Pre-COVID	January 2017–December 2019	36	3692
COVID (Stressed)	February 2020–June 2021	17	4266
Post-COVID (Train)	July 2021–October 2024	40	5047
Final Test	November 2024–October 2025	10	5078

The increase in the number of firms over time reflects listing growth and data inclusion that is survivorship neutral. Firms enter and exit the dataset dynamically based on data availability.

Regarding the daily equity prices, they are obtained from Yahoo Finance, and the adjusted open and close prices are used to compute monthly and daily returns. To calculate the returns for asset i at time t, we used the formula as follows:

$$r_{i,t} = \frac{P_{t,i}^{Close}}{P_{t,i}^{Open}} - 1 \quad (2)$$

where $P_{t,i}$ denotes the adjusted close or open price of asset i at time t, with t referring to the current day for the daily return or the current month for the monthly return, with open being the first day of the month and close the last day of the month.

The last two components are the risk-free rate, which is derived from the 13-week U.S. Treasury Bill yield (^IRX), and the S&P 500 index (^GSPC); both of them are obtained from Yahoo Finance. Risk free is annualized to convert it to daily (or monthly) risk-free rate, we applied this approximation formula:

$$r_{f,d} = \frac{r_{f,d}}{n} \quad (3)$$

with n either 252 for daily or 12 for monthly.

Overall, the data construction ensures chronological consistency: accounting variables are transformed into TTM measures, aligned to reported dates, carried forward monthly, and paired with strictly forward returns. The data structure and the financial constraints preserve investor realism and support unbiased cross-sectional learning for portfolio implementation.

4. Portfolio Construction and Allocation

The following section presents the portfolio formation as a two-step mapping from fundamentals to holdings. The first step is the stock selection: fundamentals are converted to scores, scores to ranks, and ranks to a Top-K set. The second step is weighting optimization: once a set of stocks is selected, capital is allocated either by a simple rank-based rule (monthly) or by a daily optimizer (dynamic allocation section). In the next subsections, we go through every method applied and each step.

4.1. Monthly Cross-Sectional Scaling (Rank Mapping) and Scoring Function

Fundamental variables differ in their units and distributions and can shift over time. The methodology therefore maps each feature into a cross-sectional rank within each month-end t. This produces a dimensionless signal that measures a firm's position relative to peers at that month-end. Let $x_{i,t,j}$ be feature j for stock i at month-end t. Let $rank_t(x_{i,t,j})$ be the cross-sectional rank among stocks available at t and the transformed value $\widetilde{x}_{i,t,j}$ is a rank mapped to [0, 1]:

$$\widetilde{x}_{i,t,j} = \frac{rank_t(x_{i,t,j}) - 1}{\max rank_t(x_{t,j})} \quad (4)$$

The model is trained and evaluated using $x_{i,t,j}$, not the raw features. This makes the learning problem inherently cross-sectional and directly aligned with "picking the best stocks each month." The model is a monthly ranking rule. It produces a score $s_{i,t}$ for each firm i at each month-end t. Firms are ordered by $s_{i,t}$. The score is a linear aggregation of normalized features with missing-data correction, and the feature weights are learned using a genetic algorithm. Let w_j denote the weight on feature j. Let $\widetilde{x}_{i,t,j}$ denote the month-t rank-normalized feature. Some firms may have missing values for some features. The methodology does not impute missing fundamentals; instead, it rescales the linear score by

the total absolute weight of the present features. Let $M_{i,t}$ be the set of features observed (non-missing) for firm i at month t. The score is as stated in Equation (5):

$$s_{i,t} = \frac{\sum_{j \in M_{i,t}} w_j \tilde{x}_{i,t,j}}{\sum_{j \in M_{i,t}} |w_j| + \varepsilon} \quad (5)$$

with ε a small stabilizer. The weights are normalized so that $\sum_{j=0}^{N} w_j = 1$, which makes the scale of $s_{i,t}$ comparable across candidate solutions and improves interpretability.

This score-to-ranking mapping is identical across both learning methodologies; the difference lies in the fitness function and the target variable used to optimize the chosen w.

To provide economic interpretability in the accounting variables which are used in the model, they can be grouped into four categories: profitability (e.g., ROA, operating margin), investment and growth (e.g., asset growth, sales growth), leverage (e.g., debt-to-assets), and liquidity (e.g., current ratio). These categories are well established in the asset pricing literature as predictors of cross-sectional returns.

Rather than imposing these relationships a priori, the genetic algorithm learns the relative importance of these features endogenously through optimization, allowing the model to adaptively combine signals based on their predictive contribution.

4.2. GA Ranking Objectives NDCG vs. IC

Learning-to-rank methods originated in information retrieval and have recently been applied to financial cross-sections (Poh et al., 2021; Sakagawa & Mori, 2022). Among listwise ranking metrics, Normalized Discounted Cumulative Gain (NDCG) places greater emphasis on correct ordering of highly relevant observations (Järvelin & Kekäläinen, 2002).

The genetic algorithm is implemented with a population size of 80 and evolves over 40 generations. Elite selection retains the top 15% of candidates, while mutation is applied with probability 30% and Gaussian perturbations. All candidate weight vectors are normalized to ensure comparability across solutions.

To reduce overfitting, the fitness function is computed as an average across multiple months, rather than relying on single-period performance. This temporal aggregation improves stability and reduces sensitivity to noise in individual cross-sections.

In our study, the first objective is the NDCG method, which directly rewards placing the highest-return stocks near the top of the ranking. Each month-end t, realized forward returns $r_{i,t+1}$ are converted into 12 relevance groups. This is done by ranking the realized $r_{i,t+1}$ cross-sectionally and cutting into 12 equal-count bins (duodeciles). Let $R_{i,t+1} \in \{1, 2, \ldots, 11, 12\}$ be the resulting relevance label, where 1 corresponds to the highest forward-return duodecile and, respectively, 12 to the lowest forward-return in the specific month.

Given candidate weights w, compute scores $s_{i,t}$, and sort stocks by decreasing score. Consider the top K stocks in that sorted list. The discounted cumulative gain at cutoff K is:

$$DCG_t(K) = \sum_{i=1}^{k} \frac{2^{rel_{i,t+1}} - 1}{log_2(i+1)} \quad (6)$$

And the normalized metric is:

$$NDCG_t(\mathrm{K}) = \frac{DCG_t(K)}{IDCG_t(K)} = \frac{1}{\mathrm{IDCG}_t(K)} \sum_{i=1}^{k} \frac{2^{rel_{i,t+1}} - 1}{log_2(i+1)} \quad (7)$$

with $rel_{i,t+1}$ the relevance $R_{i,t}$ of the stock at position i and time $t+1$, and $IDCG_k$ is the Ideal DCG_k, which represents the DCG_k k-value for a perfectly sorted list. Thus, $NDCG$ measures how closely the predicted ranking matches the optimal ordering and is particularly useful for optimizing portfolio returns under rank-based constraints (Liu, 2009). These

models are suitable for ranking equities, optimizing asset allocations, and prioritizing investment strategies. This objective directly aligns training with Top-K portfolio construction, reflecting the insight that investors primarily care about correct ordering at the top of the distribution (Pobrotyn & Bialobrzeski, 2021). Empirical evidence suggests that ranking-based objectives can materially improve Sharpe ratios relative to regression-based approaches (Poh et al., 2021). Financial studies have shown that such rank-based models can outperform traditional sorting methods, such as quantile portfolios or signal-based screening (Gu et al., 2020).

The fitness of w is the average across training months:

$$Fitness_{\mathrm{NDCG}}(w) = \frac{1}{T}\sum_{t \in T_{train}} NDCG_t(K) \tag{8}$$

NDCG strongly rewards correct ordering at the top (small i), which can be beneficial for top-K investing, but can also increase sensitivity to noise, regime change, or label instability if the "best" names are hard to identify consistently. In our case, K is equal to 30 in the main specification. This objective is intentionally top-heavy: it places much more value on getting the first positions correct than on globally correlating scores with returns.

The IC is widely used in quantitative asset management to evaluate stock-picking models (Grinold & Kahn, 2000). Unlike NDCG, which emphasizes top-tail ordering, IC rewards consistent monotonic alignment across the entire cross-section. Recent machine learning studies (Li, 2023) confirm cardinalilty is imporant to the stock selection and IC remains a robust evaluation metric for financial ranking models. This second objective targets cross-sectional association between scores and realized forward returns across the entire universe, not just the top K. IC measures whether higher model scores correspond to higher realized forward returns in the cross-section (Mondal et al., 2019). Each month-end t, compute the Spearman rank correlation between scores $s_{i,t}$ and realized forward returns $r_{i,t+1}$. Let ρ_t denote this correlation as follow:

$$\rho_t = Spearman\ corr(s_{i,t}, r_{i,t+1}) \tag{9}$$

Respectfully, the fitness is the average over training months:

$$Fitness_{\mathrm{IC}}(w) = \frac{1}{T}\sum_{t \in T_{train}} \rho_t \tag{10}$$

This objective rewards monotone alignment across the full cross-section, which can produce a different set of learned weights than a top-K objective, even though the underlying features and scoring rule remain unchanged. IC rewards a linear association between scores and forward returns. A model can achieve decent IC even if the exact ordering of the very top names is imperfect, as long as the overall cross-sectional relationship is positive.

Both objectives use the same GA mechanism to search for w. Each candidate w is evaluated by generating month-by-month rankings via $s_{i,t}$, then scoring those rankings using either $NDCG_t$(K) or ρ_t.

Conceptually, the GA repeats the following loop: initialize a population of weight vectors, compute fitness for each, keep the top-performing candidates ("elites"), generate new candidates by combining pairs of parents (a convex mixing of weights), randomly perturb some candidates (mutation), renormalize weights by $\sum_{j=0}^{N} w_j = 1$, and iterate for a fixed number of generations. The final w is the highest-fitness candidate found.

The output of this stage is a fixed weight vector or, depending on the objective. This w^{IC} or w^{NDCG} is then held constant in all regime evaluations and the final out-of-sample test.

The portfolio is constructed either as a top-k or as a diversified variant. In the second case, the selection is not limited to the top-ranked-performing stocks but also proceeds sequentially down the score-ordered list with a per-sector acceptance rule. In this case, the stocks are selected by selecting the best-scoring stocks per sector, with no more than 2 stocks per sector. This produces a Top-K set that respects the model's ranking while preventing sectoral concentration. Additionally, this forces diversification within Top-K and prevents the portfolio from being dominated by a single sector during sector-led markets. Once the Top-K set is chosen, we applied equal weights (1/K) or exponential weights with larger weights to higher-ranked stocks via a monotone rank-decay schedule so that the first ranks receive the greatest allocation. This is consistent with the GA-NDCG objective, which emphasizes the top of the ranking.

4.3. Dynamic Portfolio Optimization

After the first-stage stock selection, a second-stage dynamic allocation procedure is applied. The purpose of this stage is not to re-select securities, but to optimally allocate capital within the already selected universe using short-horizon risk estimates derived from daily returns. It is important to emphasize that the dynamic allocation stage does not introduce additional predictive signals or market timing. Instead, it operates purely as a conditional risk management layer. The set of selected assets remains fixed within each month, and only portfolio weights are adjusted based on rolling estimates of expected returns and covariance. This separation allows clear attribution between cross-sectional selection and risk allocation.

Let $U_m = \{1, \ldots, N_m\}$ denote the set of N_m stocks selected at the beginning of month m, either through the top-k ranking or the sector-capped selection rule. This universe remains fixed during all trading days d that belong to month m. Only portfolio weights are adjusted daily; the asset set itself does not change intra-month.

For each trading day d, define a rolling lookback window of length L trading days (in our implementation, $L = 30$). Let $r_{i,t}$ denote the realized daily return of stock i $\in U_m$ on day t. Over the window $t = d - L, \ldots, d - 1$, the conditional mean vector and covariance matrix are estimated as:

$$\mu_d = \frac{1}{L}\sum_{t=d-L}^{d-1} r_t \tag{11}$$

$$\text{and } \Sigma_d = \frac{1}{L}\sum_{t=d-L}^{d-1} (r_t - \mu_d)(r_t - \mu_d)^T \tag{12}$$

where r_t is the $N_m \times 1$ vector of asset returns at time t. These estimators provide short-horizon forecasts of expected return and risk for the current universe.

The portfolio weight vector at day d is denoted as $w_d = (w_{1,d}, \ldots..w_{N_m,d})^T$.

All allocations satisfy full-investment and long-only constraints:

$$\sum_{i=0}^{N_m} w_{i,d} = 1 \text{ and } w_{i,d} \geq 0 \tag{13}$$

These constraints ensure capital is fully deployed and short-selling is not allowed.

The first allocation method minimizes conditional variance while incorporating stability controls. The objective function is:

$$\min_{w_d} w_d^T \Sigma_d w_d + \lambda_2 \|w_d\|_2^2 + \lambda_{to} \|w_d - w_{d-1}\|_2^2 \tag{14}$$

subject to the constraints above (13). The first term, $w_d^T \Sigma_d w_d$, represents the conditional portfolio variance. The second term, $\lambda_2 \|w_d\|_2^2$, is an $\ell 2$-regularisation penalty that shrinks weights toward zero and discourages extreme concentration in a small number of stocks. The third term, $\lambda_{to} \|w_d - w_{d-1}\|_2^2$, penalizes deviations from the previous day's allocation

and therefore controls turnover. The parameters $\lambda_2 > 0$ and $\lambda_{to} > 0$ govern the strength of these stabilization effects. In simple terms, the λ_2 term discourages extreme concentration and the λ_{to} term discourages large day-to-day changes in weights.

Transaction costs are explicitly incorporated through a turnover-based model. A proportional cost of approximately 5 basis points per unit of turnover is applied, where turnover is defined as the absolute change in portfolio weights between consecutive trading days. Additionally, the optimization includes regularization terms to improve stability. An $\ell 2$ penalty ($\lambda_1 = 0.005$) discourages extreme concentration, while a turnover penalty ($\lambda_2 = 0.05$) limits excessive trading. These parameters are selected to balance stability and responsiveness and are consistent with practical portfolio implementation.

The second allocation method maximizes the conditional Sharpe ratio. Let $r_{f,d}$ denote the daily risk-free rate. The excess expected return of the portfolio is $w_d^T \mu_d - r_{f,d}$, and the conditional volatility is $\sqrt{w_d^T \Sigma_d w_d}$. The optimization problem is:

$$\max_{w_d} \frac{w_d^T \mu_d - r_{f,d}}{\sqrt{w_d^T \Sigma_d w_d}} - \lambda_2 \|w_d\|_2^2 - \lambda_{to} \|w_d - w_{d-1}\|_2^2 \tag{15}$$

subject to the same full-investment and long-only constraints.

The first term represents the conditional Sharpe ratio, i.e., expected excess return per unit of volatility. The regularization and turnover penalties serve the same purpose as in the minimum-variance formulation: reducing concentration and limiting instability across trading days.

Both optimizations in the daily portfolio include the trading costs which are applied through a turnover-based cost proportional to weight movement and the total cost is as equal to $\sum_i \left\| w_{i,d} - w_{i,d-1} \right\|$. Transaction cost modeling follows the importance of incorporating trading frictions documented in recent portfolio selection research (Li, 2023; Qiao et al., 2023).

In summary, the dynamic allocation stage operates conditionally on a fixed monthly-selected universe. Each trading day, expected returns, and covariances are estimated from rolling daily data, and weights are updated through either a penalized minimum-variance or penalized maximum-Sharpe optimization. The inclusion of regularization and turnover penalties ensures numerical stability, reduces excessive concentration, and produces implementable portfolios under realistic trading frictions.

5. Empirical Results

This section presents the empirical findings in two stages. First, we evaluate the machine learning performance of the ranking models using cross-sectional metrics. Second, we assess the financial performance of the resulting portfolios at both the monthly and daily frequencies. The results are organized to clearly distinguish predictive quality from realized investment outcomes.

To enhance interpretability, we first examine the features that receive the highest and lowest weights from the genetic algorithm under both optimization objectives. Table 2 reports the top and bottom weighted accounting features learned by the GA under the NDCG and IC objectives (all the features reported in Appendix A Tables A4 and A5).

Table 2. Top & bottom weighted accounting features learned by the GA under the NDCG and IC objectives.

MODEL	TOP POSITIVE FEATURES	TOP NEGATIVE FEATURES
GA-NDCG	goodWill	noncontrollingInterestInConsolidatedEntity
	commonStockSharesOutstanding	netInvestedCapital
	stockBasedCompensation	Debt_to_Assets
	EBITDA_to_Assets	NetDebt_to_Assets
	otherCurrentAssets	netIncome_TTM
	otherLiab	ebitda_TTM
	warrants	netDebt
GA-IC	cashFlowsOtherOperating	otherStockholderEquity
	Operating_Margin	costOfRevenue
	netIncome_cf	netIncomeApplicableToCommonShares
	reconciledDepreciation	nonRecurring
	nonCurrentLiabilitiesTotal	extraordinaryItems
	ROA	Asset_Growth
	capitalSurpluse	longTermDebt

The results show that the GA-IC model places strong emphasis on profitability-related variables such as Operating Margin and ROA, as well as equity-based measures. This is consistent with established financial theory, where more profitable and financially stable firms tend to generate higher expected returns.

In contrast, the GA-NDCG model exhibits a more diverse set of influential features, including balance sheet components such as goodwill, stock-based compensation, and various asset and liability categories. This reflects the objective's focus on identifying extreme outperformers rather than maintaining consistent monotonic relationships across the entire cross-section.

Across both models, leverage-related variables such as Debt-to-Assets and Net Debt-to-Assets receive negative weights, indicating that highly leveraged firms are penalized in the ranking. Overall, the results demonstrate that the genetic algorithm learns economically meaningful relationships while maintaining flexibility in combining multiple accounting signals.

These findings provide additional interpretability, confirming that the learned ranking functions are not purely data-driven artifacts but reflect economically meaningful relationships.

5.1. Machine Learning Performance

Although the genetic algorithm searches over a large space of ranking functions, overfitting is mitigated through several design choices. First, the fitness function is averaged across multiple months, reducing sensitivity to individual observations. Second, the model is trained exclusively on a fixed historical window and evaluated on a strictly held-out test period without re-estimation. These constraints ensure that the reported results reflect genuine out-of-sample generalization rather than in-sample optimization.

The first stage of our analysis is to evaluate the cross-sectional predictive quality of the learned ranking functions independently of portfolio construction. No weighting scheme or optimization is applied at this stage. The objective is to determine whether the genetic algorithm extracts persistent information from the transformed accounting fundamentals after strict reported-date alignment and trailing-twelve-month aggregation.

Two learning objectives are implemented. The GA-NDCG model maximizes average monthly NDCG@30, thereby directly emphasizing the correct identification of the highest-return stocks. The GA-IC model maximizes average monthly Spearman rank correlation between predicted scores and realized forward returns, thereby promoting monotonic

ordering across the entire cross-section. Importantly, both models use exactly the same transformed fundamental variables and differ only in the fitness function used during training. Any performance differential, therefore, reflects objective design rather than data construction.

Figure 1 presents the time series of monthly NDCG@30 values across the full sample period, with regimes identified as Pre-COVID, COVID, Post-COVID (training), and Final Test. The figure allows assessment of three properties: stability, resilience during structural breaks, and out-of-sample persistence. First, the time-series behavior demonstrates whether ranking quality remains systematically positive rather than episodic. Second, performance during the COVID regime provides a stress test under extreme market dispersion and volatility. Third, and most importantly, the Final Test segment evaluates whether ranking power persists when the model parameters are fixed and no re-estimation occurs. The persistence of positive NDCG values during the Final Test period indicates that the learned scoring function captures structural relationships between accounting fundamentals and future returns rather than transient correlations specific to the training window. Additionally, the GA NDCG shows higher values during the training period and is slightly lower than the rest, whereas the GA IC shows similar behavior in monthly NDCG@30. That shows the GA IC model is more stable than the GA NDCG model.

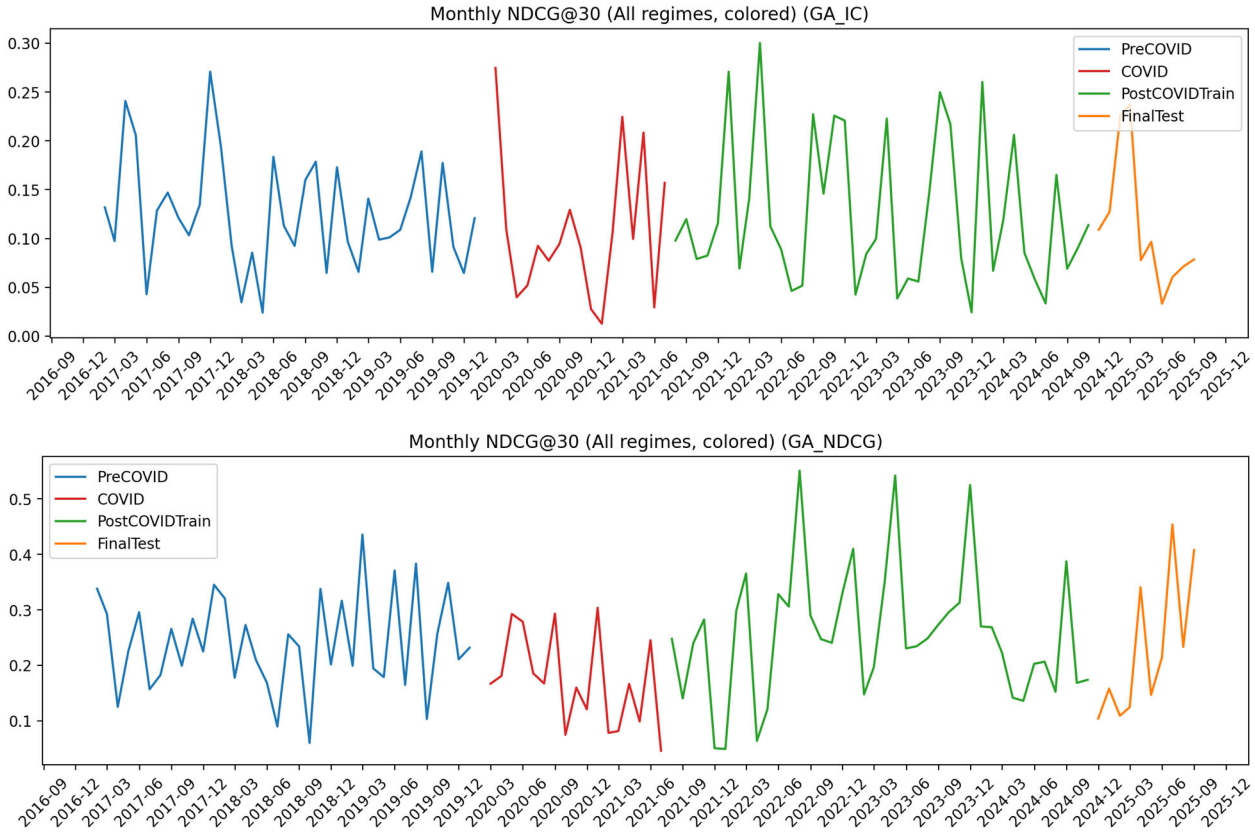


Figure 1. Monthly NDCG@30 across all regimes (GA-NDCG and GA-IC).

Table 3 summarizes average ranking quality during the training and strictly out-of-sample periods. The modest decline from 0.26 to 0.23 for GA-NDCG and from 0.13 to 0.11 for GA-IC indicates limited overfitting and stable generalization.

Table 3. Average ML metrics by regime (NDCG and IC).

Regime	Model	Average NDCG@30
Post-COVID (Train)	NDCG	0.26
	IC	0.13
Final Test	NDCG	0.23
	IC	0.11

The magnitude of NDCG is economically meaningful. NDCG is normalized such that a value of zero corresponds to random ordering and one corresponds to perfect ranking of the top-return segment. In large cross-sections with thousands of stocks, achieving an average NDCG@30 of 0.23 implies that the model systematically concentrates a disproportionate share of the highest realized returns within the top 30 predicted stocks. Because NDCG exponentially rewards correct placement of the very highest-return observations, even incremental increases represent substantial improvements in capital concentration efficiency.

In practical terms, a stable NDCG near 0.23 in an out-of-sample period implies that the probability of selecting extreme outperformers is materially above random allocation. Given the universe's dimensionality and the strict information-timing controls imposed during data construction, this level of ranking quality reflects an economically exploitable structure rather than statistical noise.

The contrast between GA-NDCG and GA-IC further clarifies the economic consequences of objective choice. Direct optimization of top-tail ranking (GA-NDCG) produces stronger concentration ability, while the IC objective emphasizes broader monotonic association across the full cross-section. Because both models use identical fundamentals, differences in performance are attributable solely to the learning criterion.

Overall, the persistence of positive NDCG and IC values in the Final Test period establishes that transformed accounting fundamentals retain cross-sectional predictive power beyond the estimation window. This statistical evidence provides the necessary foundation for evaluating whether ranking quality translates into economically significant portfolio performance in the subsequent section.

5.2. Monthly Portfolio Performance (Selection Stage)

This section evaluates whether the cross-sectional ranking signal translates into economically meaningful portfolio performance once implemented as investable monthly portfolios. Crucially, the ranking model is not re-estimated in this stage: the scoring function learned during the post-COVID training window is held fixed, and portfolios are formed mechanically each month from the resulting cross-sectional ordering.

Each month, stocks are sorted by the learned score and a Top-k subset is selected. Two selection rules are considered. The Base rule selects the highest-scoring firms without restriction, allowing the portfolio to inherit any sector clustering implied by the signal. The SectorCap2 rule constrains the selection step by limiting the portfolio to at most two stocks per sector, forcing broader industry coverage. Conditional on selection, portfolio weights are either equal-weighted or exponentially weighted, with higher-ranked names receiving systematically larger allocations. This design allows the analysis to distinguish three separate questions: (i) whether the ranking contains alpha, (ii) whether the alpha is concentrated in the very top of the ranking, and (iii) whether diversification constraints improve robustness across regimes.

In addition, portfolio robustness is assessed across multiple portfolio constructions, including different portfolio sizes, weighting schemes, and the commonly used S&P 500 benchmark. These variations serve as practical benchmarks that resemble standard fundamental investment strategies. The consistency of results across these configurations

indicates that performance is driven by the underlying ranking signal rather than specific portfolio construction choices. The investment universe consists of a broad cross-section of U.S. equities, with Top-K portfolios typically containing between 10 and 50 stocks.

Table 4 reports annualized excess return and the full set of risk and attribution statistics used throughout the study (volatility, Sharpe, Sortino, maximum drawdown, Calmar, Treynor, and alpha) for the last two regimes: PostCOVIDTrain and FinalTest (with the earlier regimes reported in Appendix A Table A2). The interpretation emphasizes out-of-sample behaviour in the FinalTest window, where the ranking model is applied without retraining.

A first result is that out-of-sample performance is strongly positive for the NDCG-trained strategies, especially when the selection step includes the sector cap. In the FinalTest regime, the GA-NDCG SectorCap2 portfolios exhibit particularly high excess returns and strong risk-adjusted performance across both weighting schemes. For example, SectorCap2 under equal weighting produces very large annualized excess returns (e.g., roughly 0.63–0.78 depending on k), while maintaining moderate drawdowns (around -0.18 to -0.22). The same pattern persists under exponential weighting, where excess returns remain high and Sharpe ratios remain well above one across k. This combination—high excess return with controlled drawdown—suggests that, in an out-of-sample environment, the signal's economic value is not limited to a narrow set of sectors and that enforcing diversification strengthens implementation rather than diluting alpha.

Second, the Base (unconstrained) GA-NDCG portfolios display a meaningful dependence on k and on the weighting rule, with a clear implication for where the signal resides. In FinalTest, Base portfolios are positive and economically relevant across k, and exponential weighting systematically improves performance relative to equal weighting at the same k. This indicates that the model's predictive information is more concentrated near the top of the ranking, consistent with the logic of optimizing the NDCG objective. In other words, when capital is tilted toward the highest-scoring stocks, realized performance improves, implying that "rank intensity" is economically informative rather than merely statistical.

Third, GA-IC portfolios also perform well out of sample, but the magnitude is smaller than the strongest GA-NDCG configurations. In FinalTest, GA-IC strategies—especially under Base selection—exhibit positive excess return and high Sharpe ratios, often exceeding one, with very small maximum drawdowns in the FinalTest window (on the order of a few percentage points). This profile is consistent with what the IC objective incentivizes: not necessarily extreme top-tail concentration, but broad monotonic alignment between scores and forward returns that can generate a smoother return process. The fact that GA-IC produces strong Sharpe with comparatively modest excess returns is economically meaningful: it suggests that the IC objective may deliver a more stable "ranking everywhere" signal, whereas the NDCG objective tends to generate more economically aggressive top-tail bets.

A fourth and important finding is the regime-dependence of the SectorCap2 constraint. In PostCOVIDTrain, SectorCap2 under GA-NDCG is often weaker (and in several configurations negative) despite strong performance in FinalTest. This pattern is informative rather than problematic. It indicates that the sector cap does not mechanically increase backtest returns; instead, it serves as a robustness device that becomes valuable when the market environment shifts. If the training period rewarded sector clustering (e.g., because the best-scoring fundamentals concentrated in a subset of industries), then removing that clustering via caps can reduce performance in-sample. However, when the environment changes out of sample, the same constraint can protect the strategy from hidden-sector bets and improve realized performance. The out-of-sample dominance of SectorCap2 is therefore consistent with the interpretation that part of the unconstrained signal may reflect time-varying sector exposures, while the constrained portfolios retain more firm-level information.

Table 4. Monthly portfolio performance by train and test regime.

Strategy	Method	Weighting	Regime	Return	Vol	Sharpe	Sortino	MaxDD	Calmar	Treynor	Alpha
NDCG@10	Base	Equal	PostCOVIDTrain	0.105	0.480	0.210	0.349	−0.588	0.242	−0.534	0.118
NDCG@20	Base	Equal	PostCOVIDTrain	0.055	0.426	0.126	0.216	−0.596	0.152	0.215	0.040
NDCG@30	Base	Equal	PostCOVIDTrain	−0.017	0.419	−0.042	−0.068	−0.620	0.026	−0.154	−0.024
NDCG@50	Base	Equal	PostCOVIDTrain	−0.073	0.375	−0.201	−0.294	−0.620	−0.066	0.411	−0.063
NDCG@10	Base	Equal	FinalTest	0.282	0.432	0.581	1.013	−0.185	1.807	−3.650	0.294
NDCG@20	Base	Equal	FinalTest	0.166	0.419	0.369	0.585	−0.220	0.975	1.884	0.154
NDCG@30	Base	Equal	FinalTest	0.154	0.397	0.362	0.547	−0.233	0.867	0.834	0.128
NDCG@50	Base	Equal	FinalTest	0.216	0.396	0.498	0.797	−0.231	1.150	0.513	0.154
NDCG@10	Base	Exponential	PostCOVIDTrain	0.180	0.480	0.348	0.600	−0.568	0.387	−0.808	0.195
NDCG@20	Base	Exponential	PostCOVIDTrain	0.116	0.430	0.256	0.432	−0.575	0.267	1.296	0.110
NDCG@30	Base	Exponential	PostCOVIDTrain	0.076	0.424	0.172	0.285	−0.585	0.191	1.559	0.073
NDCG@50	Base	Exponential	PostCOVIDTrain	0.049	0.405	0.118	0.191	−0.587	0.144	−27.940	0.049
NDCG@10	Base	Exponential	FinalTest	0.360	0.415	0.752	1.419	−0.144	2.879	−2.860	0.381
NDCG@20	Base	Exponential	FinalTest	0.268	0.408	0.587	1.004	−0.177	1.809	6.927	0.262
NDCG@30	Base	Exponential	FinalTest	0.252	0.396	0.573	0.952	−0.186	1.638	2.475	0.237
NDCG@50	Base	Exponential	FinalTest	0.254	0.392	0.584	0.969	−0.189	1.615	1.624	0.230
NDCG@10	SectorCap2	Equal	PostCOVIDTrain	−0.077	0.475	−0.169	−0.249	−0.605	−0.076	0.134	−0.047
NDCG@20	SectorCap2	Equal	PostCOVIDTrain	−0.169	0.396	−0.465	−0.603	−0.694	−0.203	0.409	−0.149
NDCG@30	SectorCap2	Equal	PostCOVIDTrain	−0.140	0.363	−0.413	−0.545	−0.629	−0.176	0.448	−0.125
NDCG@50	SectorCap2	Equal	PostCOVIDTrain	−0.140	0.363	−0.413	−0.545	−0.629	−0.176	0.448	−0.125
NDCG@10	SectorCap2	Equal	FinalTest	0.784	0.331	1.791	4.093	−0.185	4.633	0.681	0.553
NDCG@20	SectorCap2	Equal	FinalTest	0.720	0.390	1.423	3.085	−0.213	3.703	0.554	0.470
NDCG@30	SectorCap2	Equal	FinalTest	0.628	0.410	1.212	2.755	−0.224	3.102	0.448	0.373
NDCG@50	SectorCap2	Equal	FinalTest	0.628	0.410	1.212	2.755	−0.224	3.102	0.448	0.373
NDCG@10	SectorCap2	Exponential	PostCOVIDTrain	−0.013	0.482	−0.027	−0.041	−0.583	0.036	0.022	0.021
NDCG@20	SectorCap2	Exponential	PostCOVIDTrain	−0.107	0.422	−0.266	−0.366	−0.636	−0.120	0.210	−0.081
NDCG@30	SectorCap2	Exponential	PostCOVIDTrain	−0.098	0.403	−0.254	−0.350	−0.607	−0.110	0.215	−0.074
NDCG@50	SectorCap2	Exponential	PostCOVIDTrain	−0.098	0.403	−0.254	−0.350	−0.607	−0.110	0.215	−0.074
NDCG@10	SectorCap2	Exponential	FinalTest	0.714	0.310	1.779	4.493	−0.161	4.854	0.745	0.527
NDCG@20	SectorCap2	Exponential	FinalTest	0.810	0.369	1.650	4.000	−0.181	4.872	0.641	0.555
NDCG@30	SectorCap2	Exponential	FinalTest	0.758	0.377	1.533	3.739	−0.190	4.367	0.578	0.501
NDCG@50	SectorCap2	Exponential	FinalTest	0.758	0.377	1.533	3.739	−0.190	4.367	0.578	0.501
IC@10	Base	Equal	PostCOVIDTrain	0.107	0.143	0.713	1.229	−0.134	1.073	−0.971	0.114
IC@20	Base	Equal	PostCOVIDTrain	0.040	0.130	0.302	0.445	−0.154	0.486	−0.305	0.047
IC@30	Base	Equal	PostCOVIDTrain	0.053	0.133	0.386	0.606	−0.127	0.696	−0.343	0.062
IC@50	Base	Equal	PostCOVIDTrain	0.077	0.138	0.540	0.869	−0.113	1.001	−0.371	0.090
IC@10	Base	Equal	FinalTest	0.126	0.090	1.326	2.378	−0.018	9.348	6.454	0.123
IC@20	Base	Equal	FinalTest	0.119	0.087	1.299	2.560	−0.019	8.741	5.205	0.116
IC@30	Base	Equal	FinalTest	0.083	0.103	0.780	1.194	−0.030	4.265	−13.761	0.084
IC@50	Base	Equal	FinalTest	0.077	0.117	0.639	0.926	−0.032	3.863	−0.693	0.092
IC@10	Base	Exponential	PostCOVIDTrain	0.104	0.144	0.690	1.183	−0.130	1.087	−1.119	0.110
IC@20	Base	Exponential	PostCOVIDTrain	0.057	0.133	0.415	0.646	−0.145	0.638	−0.515	0.063
IC@30	Base	Exponential	PostCOVIDTrain	0.061	0.132	0.450	0.712	−0.133	0.732	−0.500	0.068
IC@50	Base	Exponential	PostCOVIDTrain	0.066	0.133	0.485	0.771	−0.131	0.782	−0.490	0.074

Table 4. *Cont.*

Strategy	Method	Weighting	Regime	Return	Vol	Sharpe	Sortino	MaxDD	Calmar	Treynor	Alpha
IC@10	Base	Exponential	FinalTest	0.143	0.096	1.398	2.635	−0.023	8.147	2.661	0.136
IC@20	Base	Exponential	FinalTest	0.118	0.084	1.337	2.471	−0.023	7.258	2.409	0.111
IC@30	Base	Exponential	FinalTest	0.105	0.090	1.118	1.849	−0.026	5.718	3.664	0.101
IC@50	Base	Exponential	FinalTest	0.097	0.094	0.987	1.574	−0.028	5.030	19.207	0.096
IC@10	SectorCap2	Equal	PostCOVIDTrain	0.030	0.131	0.223	0.347	−0.098	0.658	−0.186	0.039
IC@20	SectorCap2	Equal	PostCOVIDTrain	0.062	0.128	0.467	0.746	−0.121	0.808	−0.414	0.071
IC@30	SectorCap2	Equal	PostCOVIDTrain	0.040	0.136	0.290	0.430	−0.169	0.446	−0.284	0.049
IC@50	SectorCap2	Equal	PostCOVIDTrain	0.040	0.136	0.290	0.430	−0.169	0.446	−0.284	0.049
IC@10	SectorCap2	Equal	FinalTest	0.128	0.122	0.989	1.911	−0.047	3.705	0.751	0.104
IC@20	SectorCap2	Equal	FinalTest	0.093	0.122	0.725	1.163	−0.045	3.099	11.394	0.091
IC@30	SectorCap2	Equal	FinalTest	0.102	0.129	0.755	1.167	−0.035	4.177	−1.053	0.115
IC@50	SectorCap2	Equal	FinalTest	0.102	0.129	0.755	1.167	−0.035	4.177	−1.053	0.115
IC@10	SectorCap2	Exponential	PostCOVIDTrain	0.037	0.133	0.277	0.431	−0.107	0.675	−0.264	0.046
IC@20	SectorCap2	Exponential	PostCOVIDTrain	0.055	0.127	0.425	0.682	−0.104	0.871	−0.399	0.064
IC@30	SectorCap2	Exponential	PostCOVIDTrain	0.046	0.130	0.347	0.539	−0.129	0.632	−0.338	0.054
IC@50	SectorCap2	Exponential	PostCOVIDTrain	0.046	0.130	0.347	0.539	−0.129	0.632	−0.338	0.054
IC@10	SectorCap2	Exponential	FinalTest	0.142	0.113	1.182	2.232	−0.045	4.241	0.862	0.119
IC@20	SectorCap2	Exponential	FinalTest	0.108	0.112	0.919	1.550	−0.045	3.425	1.501	0.098
IC@30	SectorCap2	Exponential	FinalTest	0.111	0.115	0.919	1.525	−0.041	3.838	5.999	0.108
IC@50	SectorCap2	Exponential	FinalTest	0.111	0.115	0.919	1.525	−0.041	3.838	5.999	0.108

Finally, Appendix A Table A2 shows that the framework remains economically active during PreCOVID and COVID, but with different relative strengths across configurations. The COVID regime, in particular, exhibits large dispersion in outcomes across selection rules and objectives. This is expected in a stress environment where sector and factor exposures dominate cross-sectional dispersion and where the economic mapping from fundamentals to near-term returns can change abruptly. The central result is that the strategy families that perform best out of sample (notably GA-NDCG SectorCap2) also remain viable during stress, indicating that robustness is not achieved by "smoothing away" the signal.

Figure 2 presents cumulative returns of equal-weight Top-k portfolios relative to the S&P 500 during the Final Test period. This complements Table 4 by visualizing strict out-of-sample performance under equal weighting, which isolates stock selection ability without additional concentration. A persistent separation between the portfolio equity curve and the benchmark indicates that the score-based ordering contains investable information beyond broad market exposure, while the drawdown profile provides a direct visual check on tail risk.

Figure 3 evaluates whether concentrating capital toward the highest-ranked stocks improves realized performance. Consistent with the table evidence, exponential weighting tends to raise return levels for the strongest GA-NDCG strategies without introducing commensurate deterioration in drawdowns, supporting the view that the top of the ranking drives a disproportionate share of alpha.

Figure 4 examines whether the diversification constraint changes behavior across macro regimes. The key contribution of this figure is interpretability: it shows when sector caps operate as a risk control (reducing drawdowns or volatility) versus when they reduce exposure to a profitable cluster. The evidence that SectorCap2 is most valuable in the stressed and test regimes aligns with the interpretation that sector caps improve robustness when sector composition of the winners changes.

Overall, the persistence of ranking performance reflects the economic foundation of the underlying signals. Accounting characteristics such as profitability and asset growth capture structural firm-level fundamentals rather than short-term price patterns. Prior asset-pricing research demonstrates that these variables contain cross-sectional information linked to firm investment and production decisions (Novy-Marx, 2013; Cooper et al., 2008; Green et al., 2017). Because these relationships are grounded in economic mechanisms, they are more likely to persist across regimes than purely technical signals. Lastly, Top-focused objectives naturally produce concentrated portfolios. While concentration can enhance returns when signals are strong and stable, it also increases exposure to idiosyncratic and sector-specific shocks. Therefore, strong in-sample performance under NDCG optimization may partially reflect aggressive exposure to firms exhibiting extreme accounting characteristics during the recovery period.

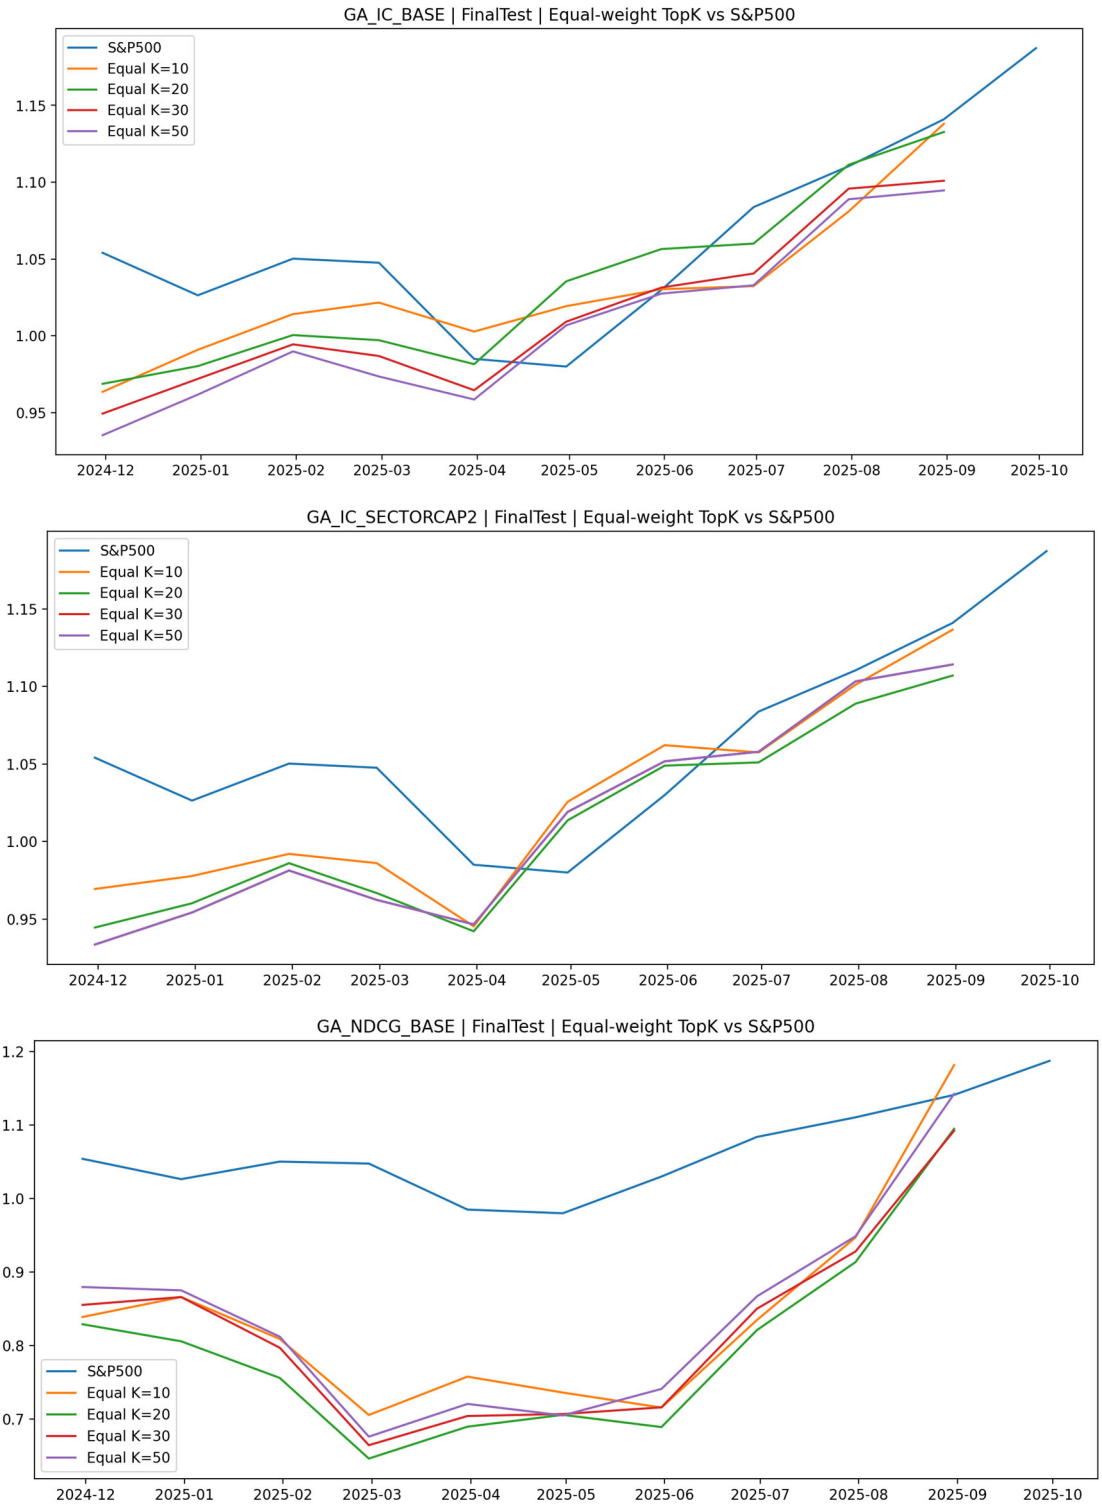


Figure 2. *Cont.*

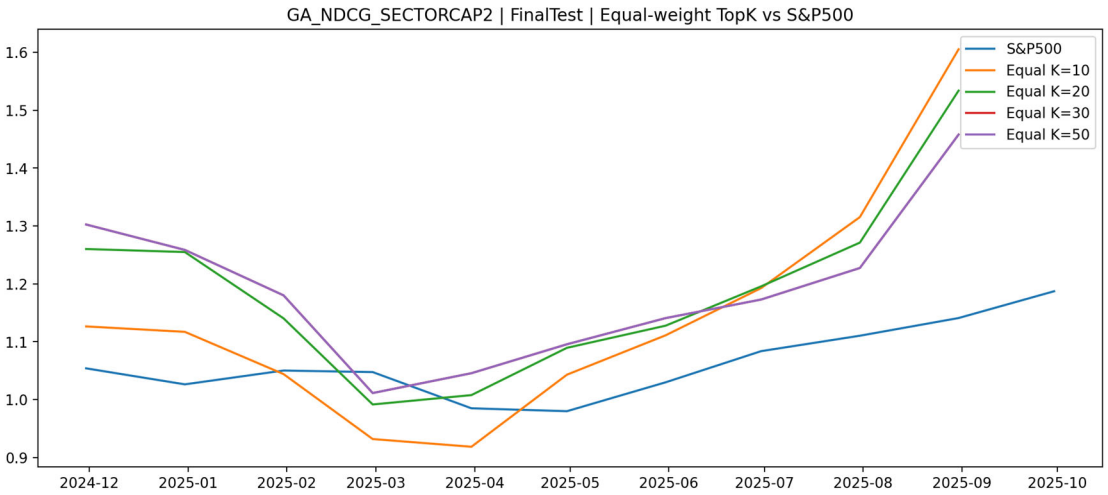


Figure 2. Final test equity curves—equal weight portfolios.

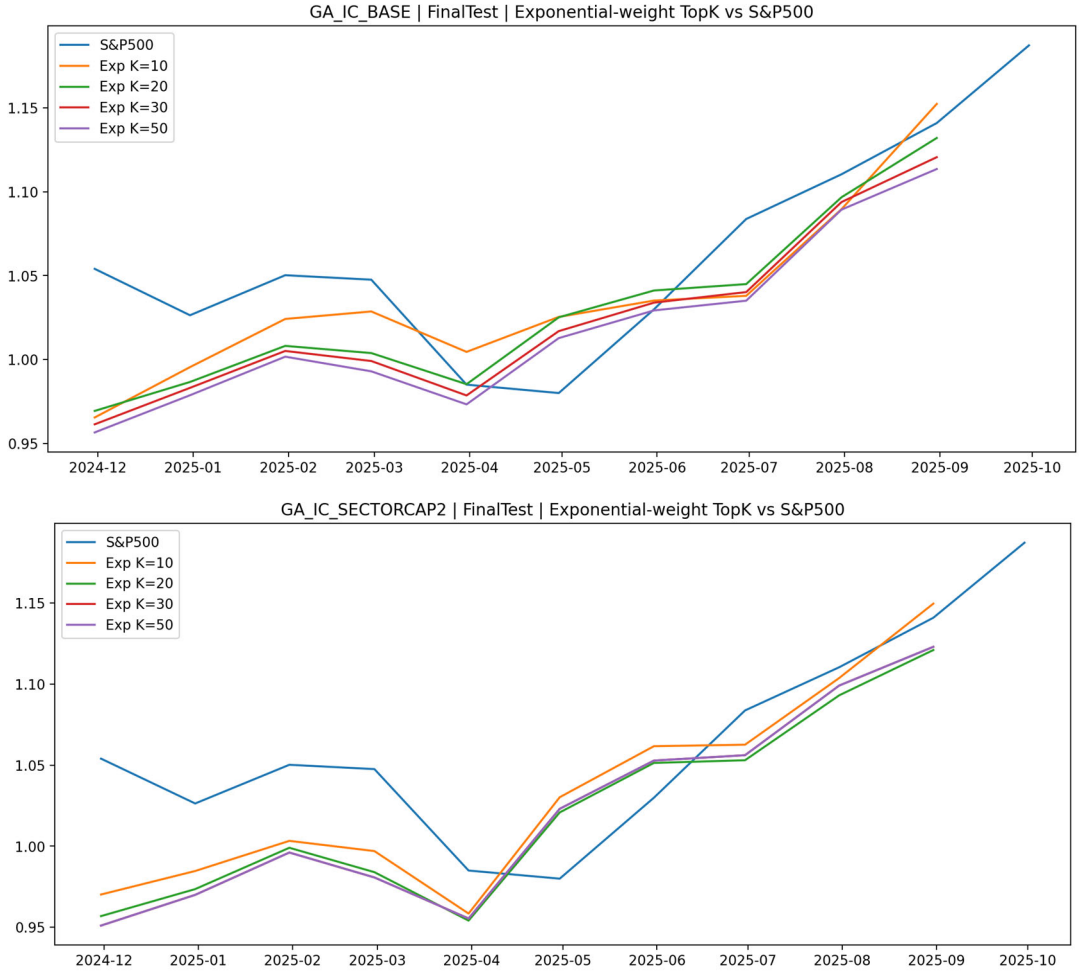


Figure 3. *Cont.*



Figure 3. Final test equity curves—exponential weight portfolios.

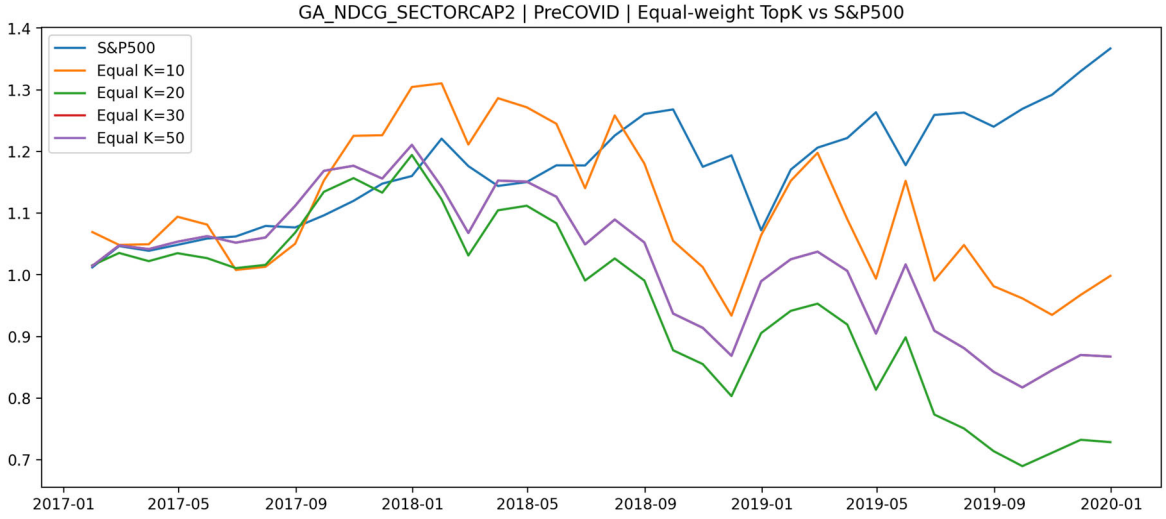


Figure 4. *Cont.*

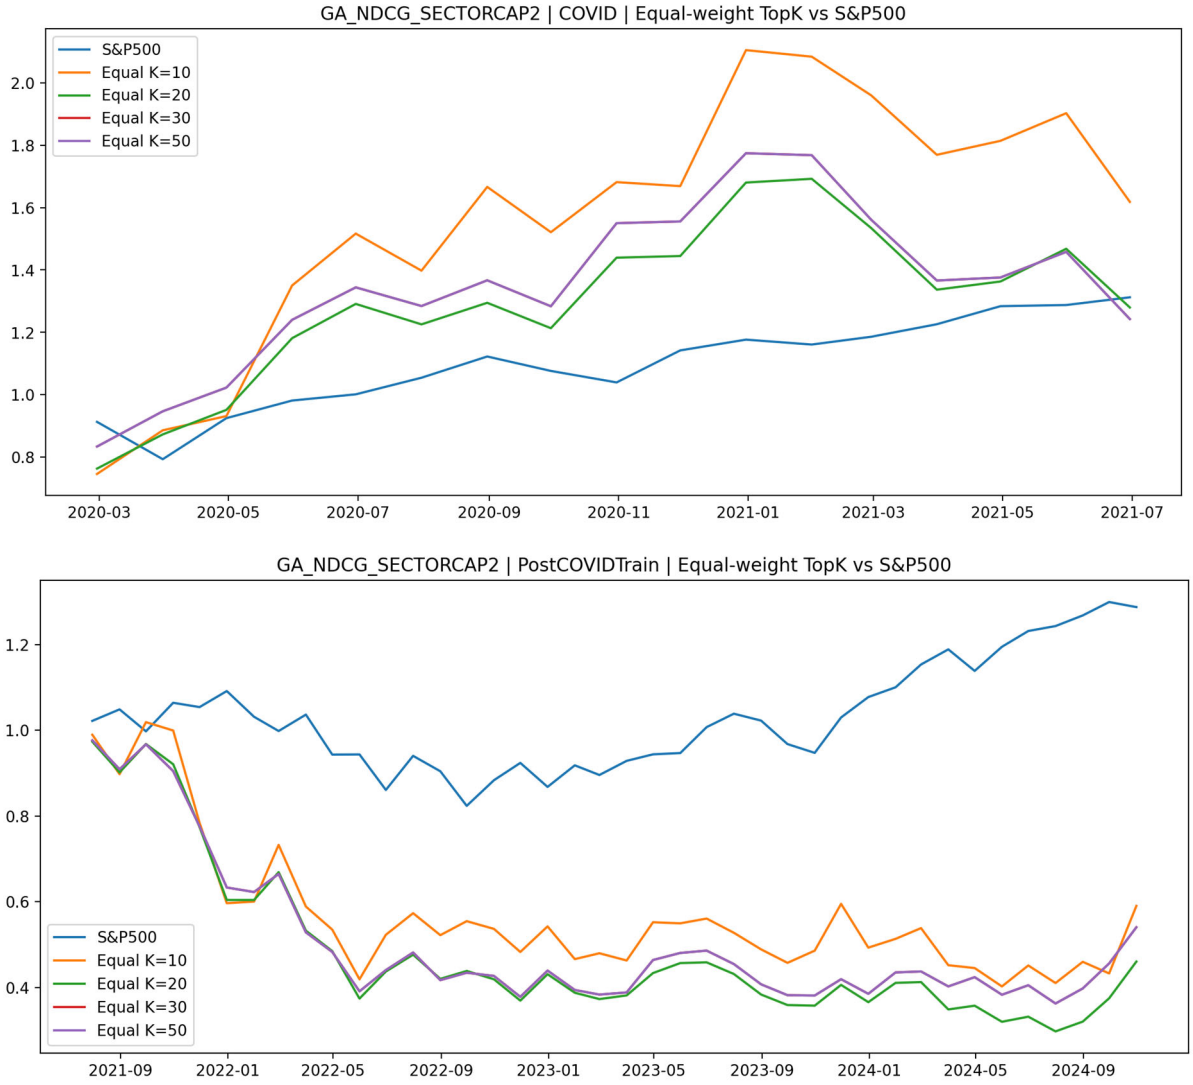


Figure 4. SectorCap2 performance across pre-COVID, COVID, and post-COVID regimes.

5.3. Daily Dynamic Allocation Results

This section evaluates whether a second-stage dynamic allocation step improves risk-adjusted performance after the monthly selection stage has fixed the investable universe. The daily optimizer does not select new stocks; it only reweights positions within the Top-50 universe determined at the monthly frequency. Hence, this stage isolates the incremental value of short-horizon risk estimation and stabilization relative to static monthly weighting.

Table 5 reports the same annualized performance statistics used in Table 4 for the dynamic allocation portfolios in PostCOVIDTrain and FinalTest (COVID results are summarized separately in Appendix A Table A3). Two allocation methods are compared: a maximum-Sharpe optimizer and a minimum-variance optimizer, both implemented with the turnover-based transaction cost model already embedded in your backtest. No management fees are applied.

Three results stand out. First, the incremental value of daily optimization is highly universe-dependent. In FinalTest, the strongest outcomes are observed when daily optimization is applied to the GA-NDCG SectorCap2 universe. Both the MaxSharpe and MinVar variants generate large positive excess returns with strong risk-adjusted statistics and moderate drawdowns. This indicates that when the selected universe is already diversified and signal-consistent, daily risk modeling can translate that selection advantage into a more efficient path of realized returns.

Table 5. Daily optimizer summary by train and test regime.

Portfolio	Optimizer	Regime	Return	Vol	Sharpe	Sortino	MaxDD	Calmar	Treynor	Alpha
GA_IC_BASE	MaxSharpe	FinalTest	−0.085	0.129	−0.687	−0.843	−0.200	−0.226	−0.319	−0.123
GA_IC_BASE	MaxSharpe	PostCOVIDTrain	0.035	0.115	0.299	0.422	−0.153	0.463	0.111	0.009
GA_IC_BASE	MinVar	FinalTest	0.067	0.124	0.524	0.726	−0.121	0.934	0.138	−0.015
GA_IC_BASE	MinVar	PostCOVIDTrain	0.108	0.127	0.804	1.154	−0.125	1.163	0.181	0.056
GA_IC_SECTORCAP2	MaxSharpe	FinalTest	0.032	0.165	0.188	0.248	−0.218	0.348	0.084	−0.030
GA_IC_SECTORCAP2	MaxSharpe	PostCOVIDTrain	0.005	0.130	0.036	0.050	−0.157	0.249	0.013	−0.024
GA_IC_SECTORCAP2	MinVar	FinalTest	0.057	0.147	0.375	0.524	−0.141	0.721	0.101	−0.037
GA_IC_SECTORCAP2	MinVar	PostCOVIDTrain	0.086	0.136	0.604	0.860	−0.159	0.773	0.129	0.029
GA_NDCG_BASE	MaxSharpe	FinalTest	−0.323	0.586	−0.665	−0.855	−0.587	−0.501	−0.425	−0.403
GA_NDCG_BASE	MaxSharpe	PostCOVIDTrain	−0.458	0.446	−1.374	−1.791	−0.897	−0.490	−0.486	−0.498
GA_NDCG_BASE	MinVar	FinalTest	0.213	0.337	0.573	0.806	−0.434	0.611	0.263	0.062
GA_NDCG_BASE	MinVar	PostCOVIDTrain	0.058	0.502	0.113	0.254	−0.622	0.152	0.055	−0.027
GA_NDCG_SECTORCAP2	MaxSharpe	FinalTest	0.573	0.449	1.009	1.965	−0.225	2.840	3.770	0.534
GA_NDCG_SECTORCAP2	MaxSharpe	PostCOVIDTrain	−0.224	0.228	−1.115	−1.472	−0.712	−0.277	−0.980	−0.238
GA_NDCG_SECTORCAP2	MinVar	FinalTest	0.504	0.256	1.597	2.577	−0.218	2.610	1.036	0.389
GA_NDCG_SECTORCAP2	MinVar	PostCOVIDTrain	−0.033	0.287	−0.117	−0.169	−0.497	0.000	−0.038	−0.098

Second, the evidence suggests that minimum variance is the more reliable daily allocator in this implementation. Several MaxSharpe configurations are unstable (including negative excess return in FinalTest for GA-IC Base and very poor outcomes for GA-NDCG Base), while MinVar is more consistently positive in both PostCOVIDTrain and FinalTest. Economically, this is a natural outcome when return estimates are noisy at short horizons: daily expected returns are difficult to estimate robustly, while covariance estimation can still provide useful risk control. The results therefore support the interpretation that the main value of daily optimization here is covariance stabilization, not short-horizon alpha forecasting.

Third, Appendix A Table A3 shows that the stressed COVID regime amplifies differences between optimizers and universes. In particular, the GA-NDCG Base universe combined with MaxSharpe performs poorly, while sector-capped universes and minimum-variance methods remain more resilient. This regime pattern reinforces the earlier conclusion from the monthly stage: unconstrained selection may embed time-varying exposures that are benign in some regimes but harmful during stress, whereas diversification constraints improve robustness.

Figure 5 provides the full time-series counterpart to Table 5, allowing visual assessment of whether improvements arise from steady compounding or from a small number of large months. The figure also clarifies whether daily optimization reduces drawdown severity relative to SPY and relative to the weakest daily strategies, which is central for evaluating implementation.

The empirical results support several important conclusions. First, the machine learning stage demonstrates that accounting fundamentals retain cross-sectional ranking power. The persistence of positive NDCG and Information Coefficient values in the Final Test period indicates that the predictive structure is not confined to the estimation window. This is especially relevant given the structural disruption introduced by COVID. The fact that ranking performance remains economically meaningful in the post-COVID environment suggests that the signal is not purely regime-specific.

Second, the difference between GA-NDCG and GA-IC highlights the importance of the objective function. Although both models use identical transformed fundamentals, optimizing directly for top-tail ranking (NDCG) tends to emphasize extreme performers, whereas optimizing rank correlation (IC) promotes monotonic consistency across the full cross-section. In financial terms, this distinction translates into differences in concentration and volatility. The empirical comparison examines whether top-focused optimization yields higher excess returns or whether broader monotonicity produces more stable portfolios.

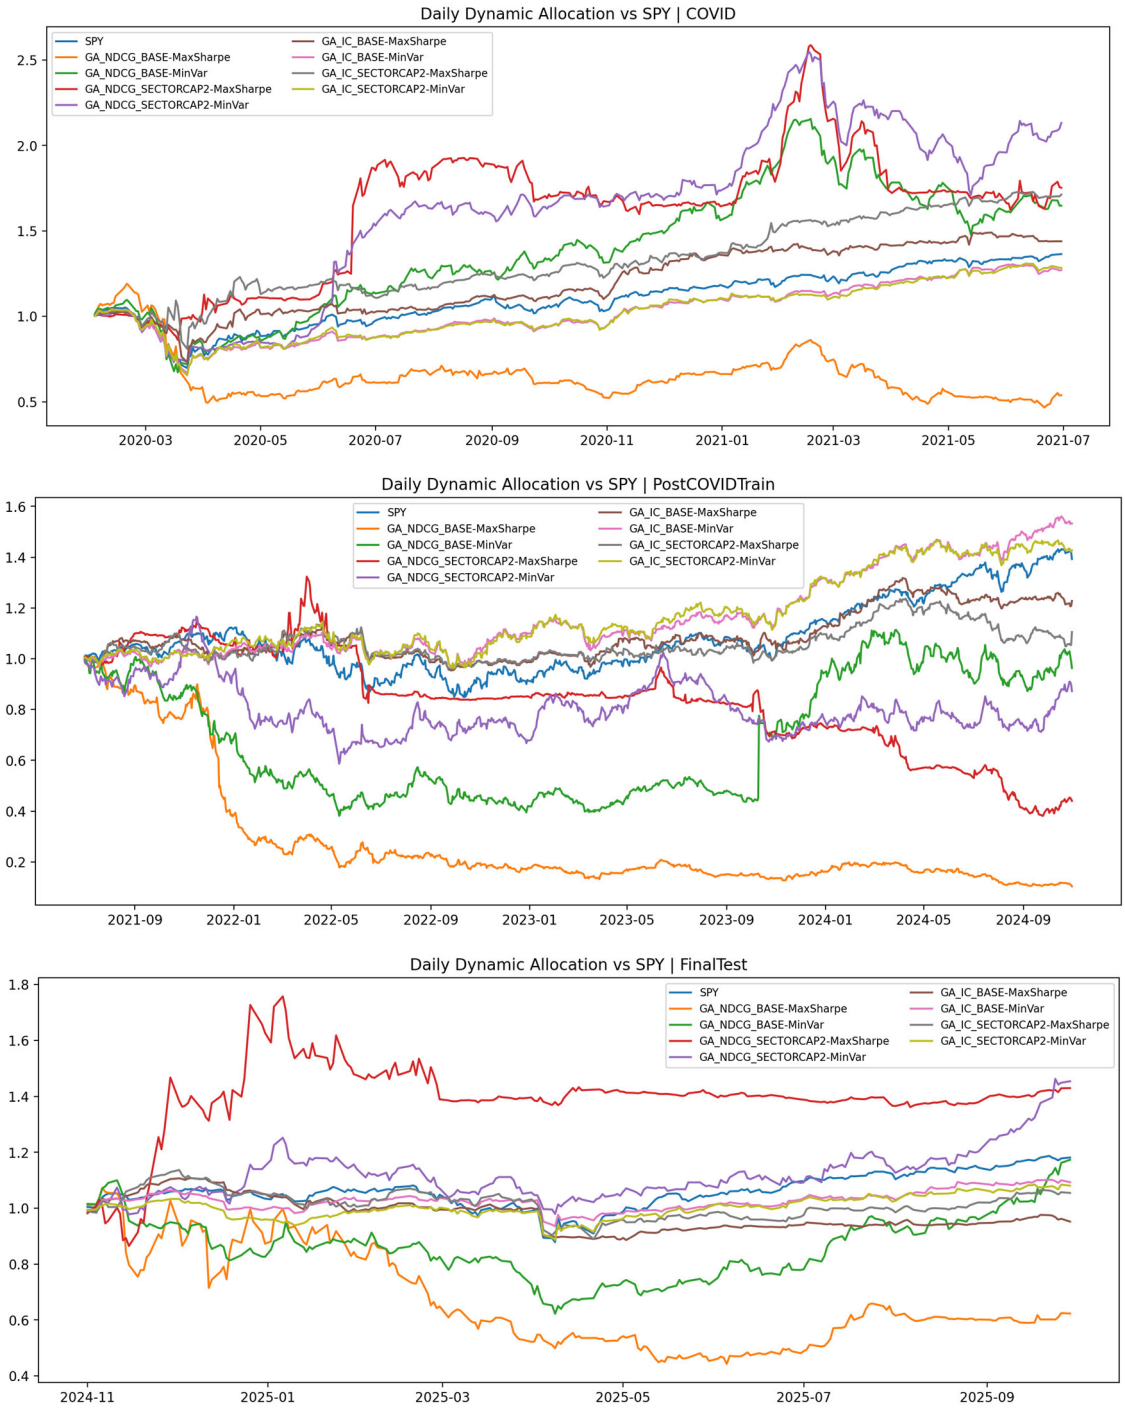


Figure 5. Daily dynamic allocation equity curves—train, test, and stressed period.

Third, the sector-capped portfolios demonstrate the trade-off between signal strength and diversification. When high-scoring firms cluster within particular industries, unconstrained portfolios may exhibit implicit sector bets. Imposing a cap of two stocks per sector reduces concentration risk. If performance remains comparable while drawdowns decline, this indicates that a portion of the raw return is driven by sector exposure rather than firm-level fundamentals. The results suggest that diversification constraints improve robustness without eliminating predictive value.

Fourth, the dynamic allocation stage contributes additional insight. Monthly selection captures cross-sectional mispricing based on fundamentals, whereas daily optimization reacts to short-horizon covariance structure. When the maximum-Sharpe or minimum-

variance allocations improve risk-adjusted performance relative to static equal-exponential weighting, this indicates that cross-sectional alpha and conditional risk modeling are complementary. Conversely, the improvements are significant for two periods, but during the training period they are marginal or nonexistent, suggesting that the greatest economic value lies in the ranking stage, which is more stable, rather than in short-term reweighting.

Importantly, the strongest evidence of robustness is observed in the Final Test period. Performance that persists strictly out-of-sample—after model estimation and without re-training—supports the economic relevance of the approach. The absence of management fees and the inclusion of turnover-based trading costs ensure that results reflect implementable net returns rather than theoretical gross performance.

Lastly, expected daily returns are noisy and difficult to estimate reliably. Small estimation errors can lead to large distortions in the weights of maximum-Sharpe portfolios. By contrast, covariance matrices are estimated using more data and are less sensitive to noise. Therefore, minimum-variance allocation provides a more stable enhancement to selection-driven alpha.

Overall, the evidence supports the view that properly timed accounting fundamentals, when optimized through ranking-consistent objectives and implemented under disciplined portfolio construction, can produce economically meaningful and reasonably robust long-only equity performance across heterogeneous market regimes.

6. Implementation and Discussion

The empirical framework developed in this study is designed not only as a methodological contribution but as an implementable investment process compatible with institutional equity management. The architecture separates cross-sectional signal extraction from conditional risk allocation, allowing each component to be evaluated independently and deployed under realistic operational constraints.

The selection stage operates monthly and relies exclusively on accounting information aligned with official filing dates and aggregated using trailing-twelve-month measures. Because the scoring function is estimated only on the designated training window and subsequently held fixed, implementation does not require continuous parameter re-estimation. Once the genetic algorithm has determined the optimal weight vector under the chosen objective, the model can be deployed mechanically: at each month-end, firms are ranked, a Top-k universe is selected under either unconstrained or sector-capped rules, and portfolio weights are assigned according to the chosen scheme. This structure closely resembles the workflow of systematic equity managers who update security selection models at discrete intervals rather than continuously.

The absence of intra-month re-selection ensures that the contribution of the ranking model can be cleanly distinguished from that of the allocation mechanism. In practice, this feature enhances transparency and governance, as it allows attribution of performance to either cross-sectional ordering or conditional risk management. Moreover, because the ranking relies on cross-sectional rank transformations rather than raw accounting magnitudes, the model remains robust to scale differences and accounting heterogeneity across firms, supporting scalability to broad equity universes.

The dynamic allocation stage, when applied, operates within the fixed monthly universe and relies on rolling daily return estimates to adjust weights subject to long-only and full-investment constraints. The inclusion of quadratic regularization and turnover penalties ensures that allocations evolve smoothly over time and that trading costs remain under control. Transaction costs are explicitly incorporated through turnover-based deductions, while management fees are excluded in order to isolate the economic contribution of the strategy itself. This design choice makes the results directly interpretable as net-of-trading but pre-fee

performance, which is standard in empirical asset pricing evaluations of implementation. Incorporating turnover penalties reduces gross returns but improves stability. This is consistent with research demonstrating that transaction costs can materially erode apparent alpha in high-turnover strategies (Li, 2023; Qiao et al., 2023; Carrasco & Kon, 2023).

From an operational perspective, the framework is computationally tractable. The genetic algorithm is trained once over the predefined training period, and daily optimization relies on standard rolling covariance estimation and constrained nonlinear programming. All required inputs—accounting data aligned to filing dates, daily prices, sector classifications, and risk-free rates—are observable in real time. The strategy, therefore, does not depend on proprietary or latency-sensitive information and can be implemented within conventional institutional infrastructure.

The empirical findings suggest that the most economically robust configuration combines objective-consistent ranking with diversification constraints. In particular, optimizing directly for top-of-ranking accuracy and imposing sector caps yields strong out-of-sample performance with controlled drawdowns. This result indicates that while the ranking model captures firm-level information embedded in fundamentals, unconstrained selection may implicitly load on time-varying sector exposures. The sector constraint mitigates this risk without eliminating the underlying signal. At the same time, the evidence indicates that daily maximum-Sharpe allocation can be sensitive to estimation noise, whereas minimum-variance allocation tends to deliver more stable improvements. This asymmetry suggests that the principal economic value arises from cross-sectional selection rather than short-horizon return forecasting.

Taken together, the framework is suitable for long-only systematic equity mandates seeking a transparent, disciplined method to incorporate accounting fundamentals into portfolio construction. Its design supports scalability across large universes and does not rely on frequent re-estimation, high-frequency signals, or unrealistic trading assumptions. The primary economic contribution lies in transforming properly timed accounting information into a stable ranking signal that remains economically meaningful out of sample.

7. Conclusions and Future Research

This study develops and evaluates a research-driven equity selection framework that integrates filing-date-aligned accounting fundamentals, learning-to-rank optimization, and constrained portfolio construction. The central motivation was to bridge the gap between documented cross-sectional return predictability in the accounting literature and the practical challenges of implementing a robust, investable long-only equity strategy.

Several structural issues were addressed. First, accounting variables were aligned to official filing dates and transformed into trailing-twelve-month measures to preserve realistic information timing. Second, stock selection was explicitly formulated as a ranking problem rather than a regression problem, aligning model optimization with the economic objective of portfolio inclusion. Third, a genetic algorithm was used to directly optimize non-differentiable ranking metrics, allowing flexible exploration of the weight space without imposing gradient-based assumptions. Finally, the selection layer was integrated with sector constraints and dynamic allocation to evaluate the interaction between cross-sectional alpha extraction and risk management.

Empirical results demonstrate that accounting-based characteristics retain economically meaningful cross-sectional ranking power. Both NDCG-optimized and IC-optimized models produce positive ranking metrics and economically significant Top-K portfolio performance. While performance declines modestly in the strictly held-out Final Test period, the persistence of positive returns suggests that the learned relationships capture structural firm-level information rather than purely transient patterns.

The comparison between ranking objectives yields several insights. Optimizing NDCG, a top-heavy metric, produces more concentrated portfolios and stronger Top-10 performance during favorable regimes. However, this concentration increases sensitivity to regime shifts and sector clustering. In contrast, optimizing the Information Coefficient promotes smoother monotonic alignment across the cross-section, leading to more stable but slightly less aggressive performance. These findings highlight that objective design materially influences portfolio concentration and risk characteristics.

Sector constraints improve out-of-sample robustness by limiting industry concentration. This suggests that part of the apparent alpha in unconstrained portfolios may reflect implicit sector exposures rather than purely firm-specific effects. Constraining sector allocation isolates stock-level fundamentals and reduces exposure to sector rotation risk.

Dynamic allocation results indicate that minimum-variance optimization provides more consistent enhancements than maximum-Sharpe allocation. This finding reinforces the well-documented difficulty of estimating short-horizon expected returns. The primary source of economic value in the framework appears to arise from cross-sectional selection rather than from high-frequency mean forecasting.

Overall, the evidence supports the conclusion that properly timed accounting fundamentals, when embedded within a ranking-consistent optimization framework and combined with disciplined portfolio construction, can generate economically meaningful and reasonably robust long-only equity performance across heterogeneous market regimes.

Several avenues for future research emerge from this study. First, hybrid objective functions combining NDCG and IC could be explored. A multi-objective genetic optimization framework may allow balancing top-tail concentration with global monotonic stability, potentially improving robustness across regimes.

Second, extending the framework to incorporate nonlinear models—such as gradient boosting or neural ranking architectures—may capture interactions between accounting variables that linear aggregation cannot fully represent. Comparing nonlinear ranking models to the interpretable linear specification would clarify the trade-off between complexity and stability.

Third, additional risk controls could be introduced, including explicit factor exposure constraints relative to Fama–French or q-factor models. This would allow clearer decomposition of alpha relative to known systematic factors.

Fourth, liquidity and capacity considerations could be incorporated to assess scalability for institutional capital. Turnover-adjusted capacity analysis would provide insight into practical implementation limits.

Fifth, adaptive or rolling re-estimation schemes could be investigated to address time variation in return predictability. While the current study enforces a strict blocked time split to ensure clean out-of-sample testing, dynamic re-training strategies may better accommodate evolving market structure.

Finally, extending the framework internationally or to alternative asset classes would test the generalizability of ranking-based accounting strategies across different institutional and informational environments. While this study focuses on U.S. equities, the proposed framework is general and can be extended to other markets. However, such extensions would require adjustments for differences in accounting standards, reporting frequency, and market liquidity.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data that support the findings of this study are not publicly available due to restrictions imposed by the data provider. The authors are not permitted to share the data. Access can be obtained directly from the data provider, subject to their licensing terms and conditions.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A

Table A1. Accounting variables and transformations (EODHD Source).

Category	Variable Name	Description	Construction
Profitability	netIncome	Net income attributable to common shareholders	Downloaded
Profitability	netIncome_TTM	Trailing 12-month net income	TTM
Profitability	ebit	Earnings before interest and taxes	Downloaded
Profitability	ebit_TTM	Trailing 12-month EBIT	TTM
Profitability	ebitda	Earnings before interest, taxes, depreciation and amortization	Downloaded
Profitability	ebitda_TTM	Trailing 12-month EBITDA	TTM
Profitability	operatingIncome	Operating income	Downloaded
Profitability	operatingIncome_TTM	Trailing 12-month operating income	TTM
Profitability	grossProfit	Gross profit	Downloaded
Profitability	grossProfit_TTM	Trailing 12-month gross profit	TTM
Profitability	Gross_Margin	Gross profit scaled by total revenue	Calculated
Profitability	Operating_Margin	Operating income scaled by total revenue	Calculated
Profitability	ROA	Net income (TTM) divided by total assets	Calculated
Profitability	netIncomeFromContinuingOps	Net income from continuing operations	Downloaded
Profitability	incomeBeforeTax	Income before taxes	Downloaded
Profitability	incomeTaxExpense	Income tax expense	Downloaded
Cash Flow	totalCashFromOperatingActivities	Operating cash flow	Downloaded
Cash Flow	totalCashFromOperatingActivities_TTM	TTM operating cash flow	TTM
Cash Flow	freeCashFlow	Free cash flow	Downloaded
Cash Flow	freeCashFlow_TTM	TTM free cash flow	TTM
Cash Flow	capitalExpenditures	Capital expenditures	Downloaded
Cash Flow	capitalExpenditures_TTM	TTM capital expenditures	TTM
Cash Flow	CFO_to_Assets	Operating cash flow (TTM)/total assets	Calculated
Cash Flow	FCF_to_Assets	Free cash flow (TTM)/total assets	Calculated
Cash Flow	depreciation	Depreciation expense	Downloaded
Cash Flow	depreciationAndAmortization	Depreciation and amortization	Downloaded
Balance Sheet	totalAssets	Total assets	Downloaded
Balance Sheet	totalLiab	Total liabilities	Downloaded
Balance Sheet	Debt_to_Assets	Total liabilities/total assets	Calculated
Balance Sheet	NetDebt_to_Assets	Net debt/total assets	Calculated
Balance Sheet	netDebt	Total debt minus cash	Downloaded

Table A1. *Cont.*

Category	Variable Name	Description	Construction
Balance Sheet	totalStockholderEquity	Shareholders' equity	Downloaded
Balance Sheet	retainedEarnings	Retained earnings	Downloaded
Balance Sheet	totalCurrentAssets	Current assets	Downloaded
Balance Sheet	totalCurrentLiabilities	Current liabilities	Downloaded
Balance Sheet	Current_Ratio	Current assets/current liabilities	Calculated
Balance Sheet	longTermDebt	Long-term debt	Downloaded
Balance Sheet	shortTermDebt	Short-term debt	Downloaded
Balance Sheet	inventory	Inventories	Downloaded
Balance Sheet	netReceivables	Net receivables	Downloaded
Balance Sheet	propertyPlantEquipment	Gross property, plant & equipment	Downloaded
Balance Sheet	propertyPlantAndEquipmentNet	Net PPE	Downloaded
Balance Sheet	goodwill	Goodwill	Downloaded
Balance Sheet	intangibleAssets	Intangible assets	Downloaded
Growth	Asset_Growth	12-month percentage change in total assets	Calculated
Growth	Sales_Growth	12-month percentage change in TTM revenue	Calculated
Revenue	totalRevenue	Quarterly total revenue	Downloaded
Revenue	totalRevenue_TTM	Trailing 12-month revenue	TTM
Expenses	costOfRevenue	Cost of revenue	Downloaded
Expenses	sellingGeneralAdministrative	SG&A expenses	Downloaded
Expenses	researchDevelopment	R&D expenses	Downloaded
Expenses	totalOperatingExpenses	Total operating expenses	Downloaded
Capital Structure	capitalStock	Issued capital stock	Downloaded
Capital Structure	additionalPaidInCapital	Additional paid-in capital	Downloaded
Capital Structure	treasuryStock	Treasury stock	Downloaded
Capital Structure	preferredStockTotalEquity	Preferred stock component	Downloaded
Other	nonOperatingIncomeNetOther	Non-operating income	Downloaded
Other	extraordinaryItems	Extraordinary items	Downloaded
Other	nonRecurring	Non-recurring items	Downloaded

Table A2. Monthly portfolio performance by pre-covid and covid regime.

Strategy	Method	Weighting	Regime	Return	Vol	Sharpe	Sortino	MaxDD	Calmar	Treynor	Alpha
NDCG@10	Base	Equal	PreCOVID	0.049	0.270	0.177	0.260	−0.420	0.157	−0.126	0.089
NDCG@20	Base	Equal	PreCOVID	0.095	0.239	0.381	0.569	−0.258	0.436	−0.412	0.119
NDCG@30	Base	Equal	PreCOVID	0.061	0.249	0.240	0.346	−0.299	0.263	−0.215	0.091
NDCG@50	Base	Equal	PreCOVID	0.060	0.253	0.230	0.321	−0.292	0.262	−0.145	0.102
NDCG@10	Base	Equal	COVID	0.218	0.534	0.373	0.622	−0.305	0.722	0.462	0.100
NDCG@20	Base	Equal	COVID	0.263	0.416	0.567	1.022	−0.276	0.961	1.142	0.202
NDCG@30	Base	Equal	COVID	0.228	0.388	0.534	0.925	−0.317	0.725	0.866	0.160
NDCG@50	Base	Equal	COVID	0.429	0.356	1.018	1.915	−0.230	1.874	2.376	0.376
NDCG@10	Base	Exponential	PreCOVID	0.056	0.262	0.210	0.305	−0.363	0.203	−0.115	0.107
NDCG@20	Base	Exponential	PreCOVID	0.087	0.235	0.357	0.512	−0.278	0.375	−0.254	0.123
NDCG@30	Base	Exponential	PreCOVID	0.078	0.236	0.318	0.456	−0.287	0.331	−0.229	0.113
NDCG@50	Base	Exponential	PreCOVID	0.078	0.237	0.320	0.455	−0.278	0.345	−0.226	0.115
NDCG@10	Base	Exponential	COVID	0.204	0.527	0.356	0.572	−0.306	0.673	0.421	0.084
NDCG@20	Base	Exponential	COVID	0.250	0.447	0.503	0.858	−0.273	0.922	0.777	0.166
NDCG@30	Base	Exponential	COVID	0.220	0.421	0.477	0.798	−0.291	0.763	0.686	0.139
NDCG@50	Base	Exponential	COVID	0.278	0.409	0.607	1.037	−0.264	1.059	0.964	0.201
NDCG@10	SectorCap2	Equal	PreCOVID	0.013	0.246	0.054	0.079	−0.287	0.104	−0.032	0.056
NDCG@20	SectorCap2	Equal	PreCOVID	−0.096	0.202	−0.499	−0.619	−0.422	−0.192	0.178	−0.047
NDCG@30	SectorCap2	Equal	PreCOVID	−0.044	0.191	−0.236	−0.325	−0.325	−0.088	0.077	0.010
NDCG@50	SectorCap2	Equal	PreCOVID	−0.044	0.191	−0.236	−0.325	−0.325	−0.088	0.077	0.010
NDCG@10	SectorCap2	Equal	COVID	0.622	0.569	0.868	1.739	−0.231	2.702	1.437	0.481
NDCG@20	SectorCap2	Equal	COVID	0.303	0.430	0.622	0.983	−0.244	1.250	1.076	0.227
NDCG@30	SectorCap2	Equal	COVID	0.258	0.399	0.581	0.947	−0.300	0.866	1.683	0.217
NDCG@50	SectorCap2	Equal	COVID	0.258	0.399	0.581	0.947	−0.300	0.866	1.683	0.217
NDCG@10	SectorCap2	Exponential	PreCOVID	−0.017	0.245	−0.072	−0.103	−0.331	−0.004	0.035	0.031
NDCG@20	SectorCap2	Exponential	PreCOVID	−0.094	0.215	−0.456	−0.576	−0.415	−0.190	0.168	−0.043
NDCG@30	SectorCap2	Exponential	PreCOVID	−0.070	0.209	−0.346	−0.456	−0.372	−0.148	0.121	−0.017
NDCG@50	SectorCap2	Exponential	PreCOVID	−0.070	0.209	−0.346	−0.456	−0.372	−0.148	0.121	−0.017
NDCG@10	SectorCap2	Exponential	COVID	0.515	0.537	0.787	1.466	−0.233	2.217	1.094	0.371
NDCG@20	SectorCap2	Exponential	COVID	0.345	0.452	0.664	1.075	−0.235	1.476	0.915	0.241
NDCG@30	SectorCap2	Exponential	COVID	0.324	0.433	0.654	1.077	−0.262	1.241	1.034	0.238
NDCG@50	SectorCap2	Exponential	COVID	0.324	0.433	0.654	1.077	−0.262	1.241	1.034	0.238
IC@10	Base	Equal	PreCOVID	0.091	0.138	0.630	1.042	−0.141	0.769	−0.220	0.135
IC@20	Base	Equal	PreCOVID	0.081	0.137	0.568	0.925	−0.190	0.517	−0.219	0.119
IC@30	Base	Equal	PreCOVID	0.070	0.138	0.490	0.759	−0.205	0.425	−0.199	0.106
IC@50	Base	Equal	PreCOVID	0.101	0.137	0.708	1.142	−0.161	0.742	−0.257	0.144
IC@10	Base	Equal	COVID	0.451	0.274	1.383	2.418	−0.036	12.493	−2.929	0.500
IC@20	Base	Equal	COVID	0.368	0.243	1.308	2.434	−0.048	7.675	−2.443	0.413
IC@30	Base	Equal	COVID	0.324	0.238	1.192	2.240	−0.049	6.713	−2.167	0.367
IC@50	Base	Equal	COVID	0.290	0.227	1.135	1.849	−0.046	6.270	−21.743	0.293

Table A2. *Cont.*

Strategy	Method	Weighting	Regime	Return	Vol	Sharpe	Sortino	MaxDD	Calmar	Treynor	Alpha
IC@10	Base	Exponential	PreCOVID	0.078	0.142	0.532	0.876	−0.156	0.614	−0.194	0.121
IC@20	Base	Exponential	PreCOVID	0.079	0.138	0.555	0.901	−0.184	0.524	−0.210	0.119
IC@30	Base	Exponential	PreCOVID	0.074	0.137	0.521	0.823	−0.193	0.472	−0.201	0.112
IC@50	Base	Exponential	PreCOVID	0.079	0.137	0.555	0.878	−0.183	0.525	−0.211	0.118
IC@10	Base	Exponential	COVID	0.412	0.278	1.258	2.028	−0.035	11.985	−3.800	0.445
IC@20	Base	Exponential	COVID	0.369	0.254	1.253	2.134	−0.044	8.402	−2.953	0.406
IC@30	Base	Exponential	COVID	0.353	0.250	1.223	2.125	−0.042	8.452	−2.721	0.390
IC@50	Base	Exponential	COVID	0.345	0.245	1.225	2.090	−0.040	8.729	−3.414	0.374
IC@10	SectorCap2	Equal	PreCOVID	0.105	0.140	0.717	1.221	−0.160	0.768	−0.254	0.149
IC@20	SectorCap2	Equal	PreCOVID	0.065	0.138	0.461	0.739	−0.203	0.408	−0.167	0.106
IC@30	SectorCap2	Equal	PreCOVID	0.082	0.128	0.614	1.020	−0.172	0.574	−0.204	0.124
IC@50	SectorCap2	Equal	PreCOVID	0.082	0.128	0.614	1.020	−0.172	0.574	−0.204	0.124
IC@10	SectorCap2	Equal	COVID	0.342	0.224	1.327	2.390	−0.027	12.903	−2.511	0.382
IC@20	SectorCap2	Equal	COVID	0.337	0.205	1.433	2.520	−0.023	14.464	−12.563	0.345
IC@30	SectorCap2	Equal	COVID	0.322	0.223	1.266	1.994	−0.024	13.734	8.249	0.311
IC@50	SectorCap2	Equal	COVID	0.322	0.223	1.266	1.994	−0.024	13.734	8.249	0.311
IC@10	SectorCap2	Exponential	PreCOVID	0.087	0.142	0.591	0.989	−0.173	0.606	−0.215	0.130
IC@20	SectorCap2	Exponential	PreCOVID	0.069	0.137	0.488	0.801	−0.196	0.440	−0.177	0.110
IC@30	SectorCap2	Exponential	PreCOVID	0.076	0.133	0.557	0.928	−0.183	0.512	−0.194	0.118
IC@50	SectorCap2	Exponential	PreCOVID	0.076	0.133	0.557	0.928	−0.183	0.512	−0.194	0.118
IC@10	SectorCap2	Exponential	COVID	0.314	0.232	1.190	1.952	−0.025	12.497	−3.123	0.343
IC@20	SectorCap2	Exponential	COVID	0.313	0.215	1.283	2.137	−0.026	11.991	−6.551	0.327
IC@30	SectorCap2	Exponential	COVID	0.309	0.221	1.228	1.963	−0.024	12.745	−18.220	0.313
IC@50	SectorCap2	Exponential	COVID	0.309	0.221	1.228	1.963	−0.024	12.745	−18.220	0.313

Table A3. Daily optimizer summary by stressed (covid) regime.

Portfolio	Optimizer	Regime	Return	Vol	Sharpe	Sortino	MaxDD	Calmar	Treynor	Alpha
GA_IC_BASE	MaxSharpe	COVID	0.335	0.257	1.125	1.768	−0.288	1.171	0.574	0.147
GA_IC_BASE	MinVar	COVID	0.219	0.245	0.809	1.070	−0.359	0.617	0.279	−0.007
GA_IC_SECTORCAP2	MaxSharpe	COVID	0.542	0.333	1.302	2.190	−0.259	2.101	0.695	0.258
GA_IC_SECTORCAP2	MinVar	COVID	0.238	0.274	0.780	1.035	−0.369	0.649	0.272	−0.015
GA_NDCG_BASE	MaxSharpe	COVID	−0.282	0.459	−0.723	−0.918	−0.608	−0.462	−0.343	−0.421
GA_NDCG_BASE	MinVar	COVID	0.521	0.365	1.150	1.557	−0.387	1.353	0.626	0.225
GA_NDCG_SECTORCAP2	MaxSharpe	COVID	0.627	0.443	1.100	2.174	−0.373	1.688	5.121	0.576
GA_NDCG_SECTORCAP2	MinVar	COVID	0.815	0.356	1.676	2.771	−0.329	2.482	1.356	0.552

Table A4. GA-IC weights (Top → Bottom).

Rank	Feature	Weight	Signal
1	cashFlowsOtherOperating	0.0237	Positive
2	Operating_Margin	0.0214	Positive
3	netIncome_cf	0.0185	Positive
4	reconciledDepreciation	0.0180	Positive
5	nonCurrentLiabilitiesTotal	0.0177	Positive
6	ROA	0.0172	Positive
7	capitalSurpluse	0.0172	Positive
8	totalStockholderEquity	0.0168	Positive
9	totalPermanentEquity	0.0163	Positive
10	goodWill	0.0154	Positive
11	liabilitiesAndStockholdersEquity	0.0151	Positive
12	netInvestedCapital	0.0147	Positive
13	Current_Ratio	0.0142	Positive
14	commonStockTotalEquity	0.0135	Positive
15	salePurchaseOfStock	0.0131	Positive
16	otherNonCashItems	0.0129	Positive
17	earningAssets	0.0129	Positive
18	preferredStockRedeemable	0.0127	Positive
19	retainedEarnings	0.0121	Positive
20	totalCashFromOperatingActivities	0.0119	Positive
21	researchDevelopment	0.0112	Positive
22	longTermInvestments	0.0106	Positive
23	grossProfit	0.0105	Positive
24	otherOperatingExpenses	0.0105	Positive
25	changeToLiabilities	0.0102	Positive
26	capitalStock	0.0099	Positive
27	changeInCash	0.0098	Positive
28	ebit	0.0095	Positive
29	totalAssets	0.0089	Positive
30	FCF_to_Assets	0.0089	Positive
31	interestIncome	0.0085	Positive
32	netReceivables	0.0082	Positive
33	propertyPlantEquipment	0.0082	Positive
34	netIncome	0.0082	Positive
35	capitalExpenditures	0.0082	Positive
36	changeReceivables	0.0077	Positive
37	freeCashFlow_TTM	0.0077	Positive
38	nonCurrrentAssetsOther	0.0073	Positive

Table A4. *Cont.*

Rank	Feature	Weight	Signal
39	totalRevenue_TTM	0.0072	Positive
40	totalLiab	0.0071	Positive
41	interestExpense	0.0069	Positive
42	incomeTaxExpense	0.0066	Positive
43	temporaryEquityRedeemableNoncontrollingInterests	0.0061	Positive
44	discontinuedOperations	0.0059	Positive
45	ebit_TTM	0.0058	Positive
46	depreciation	0.0056	Positive
47	Debt_to_Assets	0.0056	Positive
48	sellingGeneralAdministrative	0.0056	Positive
49	retainedEarningsTotalEquity	0.0055	Positive
50	noncontrollingInterestInConsolidatedEntity	0.0053	Positive
51	CFO_to_Assets	0.0052	Positive
52	ebitda	0.0051	Positive
53	netIncome_TTM	0.0051	Positive
54	cash	0.0047	Positive
55	ebitda_TTM	0.0046	Positive
56	exchangeRateChanges	0.0046	Positive
57	year	0.0046	Positive
58	incomeBeforeTax	0.0045	Positive
59	shortLongTermDebt	0.0039	Positive
60	negativeGoodwill	0.0037	Positive
61	changeToInventory	0.0032	Positive
62	shortLongTermDebtTotal	0.0032	Positive
63	capitalExpenditures_TTM	0.0030	Positive
64	minorityInterest	0.0029	Positive
65	dividendsPaid	0.0029	Positive
66	commonStock	0.0028	Positive
67	operatingIncome	0.0028	Positive
68	totalCashFromOperatingActivities_TTM	0.0028	Positive
69	Sales_Growth	0.0028	Positive
70	investments	0.0027	Positive
71	accumulatedAmortization	0.0026	Positive
72	EBITDA_to_Assets	0.0025	Positive
73	changeInWorkingCapital	0.0024	Positive
74	freeCashFlow	0.0022	Positive
75	currentDeferredRevenue	0.0019	Positive
76	delay	0.0019	Positive
77	netDebt	0.0017	Positive
78	taxProvision	0.0016	Positive

Table A4. *Cont.*

Rank	Feature	Weight	Signal
79	endPeriodCashFlow	0.0015	Positive
80	inventory	0.0014	Positive
81	totalCurrentAssets	0.0012	Positive
82	treasuryStock	0.0000	Positive
83	beginPeriodCashFlow	−0.0003	Negative
84	totalCashFromFinancingActivities	−0.0003	Negative
85	totalCashflowsFromInvestingActivities	−0.0004	Negative
86	otherAssets	−0.0005	Negative
87	netIncomeFromContinuingOps	−0.0005	Negative
88	nonCurrentAssetsTotal	−0.0006	Negative
89	longTermDebtTotal	−0.0007	Negative
90	shortTermDebt	−0.0007	Negative
91	netTangibleAssets	−0.0007	Negative
92	preferredStockTotalEquity	−0.0009	Negative
93	totalCurrentLiabilities	−0.0013	Negative
94	deferredLongTermLiab	−0.0013	Negative
95	cashAndCashEquivalentsChanges	−0.0014	Negative
96	nonCurrentLiabilitiesOther	−0.0015	Negative
97	changeToOperatingActivities	−0.0017	Negative
98	cashAndEquivalents	−0.0017	Negative
99	preferredStockAndOtherAdjustments	−0.0019	Negative
100	otherCurrentAssets	−0.0019	Negative
101	otherCashflowsFromFinancingActivities	−0.0019	Negative
102	changeToAccountReceivables	−0.0022	Negative
103	NetDebt_to_Assets	−0.0023	Negative
104	totalRevenue	−0.0025	Negative
105	otherCashflowsFromInvestingActivities	−0.0028	Negative
106	intangibleAssets	−0.0034	Negative
107	netInterestIncome	−0.0036	Negative
108	otherCurrentLiab	−0.0039	Negative
109	changeToNetincome	−0.0039	Negative
110	nonOperatingIncomeNetOther	−0.0041	Negative
111	cashAndShortTermInvestments	−0.0045	Negative
112	warrants	−0.0045	Negative
113	totalOtherIncomeExpenseNet	−0.0049	Negative
114	netBorrowings	−0.0049	Negative
115	depreciationAndAmortization	−0.0051	Negative
116	Gross_Margin	−0.0052	Negative
117	grossProfit_TTM	−0.0052	Negative
118	propertyPlantAndEquipmentNet	−0.0054	Negative

Table A4. *Cont.*

Rank	Feature	Weight	Signal
119	propertyPlantAndEquipmentGross	−0.0061	Negative
120	accumulatedOtherComprehensiveIncome	−0.0063	Negative
121	otherLiab	−0.0068	Negative
122	effectOfAccountingCharges	−0.0070	Negative
123	capitalLeaseObligations	−0.0079	Negative
124	additionalPaidInCapital	−0.0082	Negative
125	operatingIncome_TTM	−0.0085	Negative
126	commonStockSharesOutstanding	−0.0088	Negative
127	otherItems	−0.0090	Negative
128	deferredLongTermAssetCharges	−0.0091	Negative
129	stockBasedCompensation	−0.0096	Negative
130	accumulatedDepreciation	−0.0102	Negative
131	netWorkingCapital	−0.0110	Negative
132	issuanceOfCapitalStock	−0.0111	Negative
133	totalOperatingExpenses	−0.0112	Negative
134	sellingAndMarketingExpenses	−0.0116	Negative
135	accountsPayable	−0.0116	Negative
136	shortTermInvestments	−0.0116	Negative
137	longTermDebt	−0.0121	Negative
138	Asset_Growth	−0.0124	Negative
139	extraordinaryItems	−0.0137	Negative
140	nonRecurring	−0.0142	Negative
141	netIncomeApplicableToCommonShares	−0.0144	Negative
142	costOfRevenue	−0.0147	Negative
143	otherStockholderEquity	−0.0159	Negative

Table A5. GA-NDCG weights (Top → Bottom).

Rank	Feature	Weight	Signal
1	goodWill	0.0179	Positive
2	commonStockSharesOutstanding	0.0177	Positive
3	stockBasedCompensation	0.0159	Positive
4	EBITDA_to_Assets	0.0156	Positive
5	otherCurrentAssets	0.0143	Positive
6	otherLiab	0.0143	Positive
7	warrants	0.0140	Positive
8	deferredLongTermAssetCharges	0.0133	Positive
9	longTermDebtTotal	0.0120	Positive
10	dividendsPaid	0.0117	Positive
11	totalCashFromFinancingActivities	0.0116	Positive
12	netWorkingCapital	0.0113	Positive
13	otherNonCashItems	0.0112	Positive

Table A5. *Cont.*

Rank	Feature	Weight	Signal
14	otherCashflowsFromFinancingActivities	0.0110	Positive
15	otherItems	0.0108	Positive
16	interestExpense	0.0107	Positive
17	negativeGoodwill	0.0102	Positive
18	cashAndShortTermInvestments	0.0095	Positive
19	retainedEarnings	0.0093	Positive
20	totalOperatingExpenses	0.0092	Positive
21	delay	0.0092	Positive
22	Asset_Growth	0.0087	Positive
23	currentDeferredRevenue	0.0086	Positive
24	netIncomeFromContinuingOps	0.0081	Positive
25	netTangibleAssets	0.0080	Positive
26	capitalExpenditures_TTM	0.0079	Positive
27	Current_Ratio	0.0068	Positive
28	capitalStock	0.0067	Positive
29	propertyPlantAndEquipmentGross	0.0067	Positive
30	propertyPlantAndEquipmentNet	0.0065	Positive
31	nonCurrentLiabilitiesTotal	0.0055	Positive
32	investments	0.0053	Positive
33	netIncomeApplicableToCommonShares	0.0052	Positive
34	operatingIncome_TTM	0.0046	Positive
35	changeToAccountReceivables	0.0046	Positive
36	totalOtherIncomeExpenseNet	0.0045	Positive
37	treasuryStock	0.0039	Positive
38	CFO_to_Assets	0.0037	Positive
39	interestIncome	0.0037	Positive
40	sellingGeneralAdministrative	0.0034	Positive
41	otherStockholderEquity	0.0033	Positive
42	otherCurrentLiab	0.0031	Positive
43	cashAndEquivalents	0.0031	Positive
44	changeToLiabilities	0.0030	Positive
45	effectOfAccountingCharges	0.0030	Positive
46	Sales_Growth	0.0026	Positive
47	netIncome_cf	0.0026	Positive
48	costOfRevenue	0.0024	Positive
49	beginPeriodCashFlow	0.0021	Positive
50	exchangeRateChanges	0.0019	Positive
51	retainedEarningsTotalEquity	0.0017	Positive
52	endPeriodCashFlow	0.0013	Positive
53	inventory	0.0012	Positive
54	nonCurrentLiabilitiesOther	0.0011	Positive

Table A5. *Cont.*

Rank	Feature	Weight	Signal
55	changeInWorkingCapital	0.0010	Positive
56	accumulatedDepreciation	0.0007	Positive
57	ebit_TTM	0.0006	Positive
58	ebitda	0.0006	Positive
59	cashAndCashEquivalentsChanges	0.0005	Positive
60	netIncome	0.0003	Positive
61	otherOperatingExpenses	−0.0001	Negative
62	totalCurrentLiabilities	−0.0001	Negative
63	totalRevenue_TTM	−0.0002	Negative
64	additionalPaidInCapital	−0.0002	Negative
65	totalCashFromOperatingActivities_TTM	−0.0003	Negative
66	incomeBeforeTax	−0.0004	Negative
67	reconciledDepreciation	−0.0005	Negative
68	earningAssets	−0.0006	Negative
69	totalStockholderEquity	−0.0012	Negative
70	totalCurrentAssets	−0.0012	Negative
71	shortTermDebt	−0.0014	Negative
72	longTermInvestments	−0.0014	Negative
73	longTermDebt	−0.0016	Negative
74	commonStockTotalEquity	−0.0020	Negative
75	depreciationAndAmortization	−0.0021	Negative
76	capitalSurpluse	−0.0026	Negative
77	otherAssets	−0.0027	Negative
78	operatingIncome	−0.0027	Negative
79	totalAssets	−0.0028	Negative
80	changeToOperatingActivities	−0.0029	Negative
81	preferredStockAndOtherAdjustments	−0.0031	Negative
82	accountsPayable	−0.0036	Negative
83	freeCashFlow	−0.0037	Negative
84	accumulatedAmortization	−0.0037	Negative
85	otherCashflowsFromInvestingActivities	−0.0037	Negative
86	extraordinaryItems	−0.0040	Negative
87	totalCashFromOperatingActivities	−0.0042	Negative
88	grossProfit_TTM	−0.0042	Negative
89	commonStock	−0.0042	Negative
90	preferredStockTotalEquity	−0.0044	Negative
91	researchDevelopment	−0.0045	Negative
92	nonCurrrentAssetsOther	−0.0045	Negative
93	intangibleAssets	−0.0047	Negative
94	preferredStockRedeemable	−0.0049	Negative
95	capitalLeaseObligations	−0.0049	Negative

Table A5. *Cont.*

Rank	Feature	Weight	Signal
96	ROA	−0.0051	Negative
97	netInterestIncome	−0.0052	Negative
98	minorityInterest	−0.0054	Negative
99	cash	−0.0055	Negative
100	discontinuedOperations	−0.0057	Negative
101	changeToNetincome	−0.0060	Negative
102	propertyPlantEquipment	−0.0060	Negative
103	nonOperatingIncomeNetOther	−0.0061	Negative
104	changeReceivables	−0.0061	Negative
105	changeInCash	−0.0061	Negative
106	liabilitiesAndStockholdersEquity	−0.0062	Negative
107	capitalExpenditures	−0.0062	Negative
108	freeCashFlow_TTM	−0.0063	Negative
109	changeToInventory	−0.0064	Negative
110	nonRecurring	−0.0066	Negative
111	shortTermInvestments	−0.0066	Negative
112	deferredLongTermLiab	−0.0066	Negative
113	Gross_Margin	−0.0074	Negative
114	shortLongTermDebt	−0.0075	Negative
115	cashFlowsOtherOperating	−0.0076	Negative
116	accumulatedOtherComprehensiveIncome	−0.0078	Negative
117	netReceivables	−0.0081	Negative
118	totalLiab	−0.0083	Negative
119	totalCashflowsFromInvestingActivities	−0.0083	Negative
120	issuanceOfCapitalStock	−0.0095	Negative
121	temporaryEquityRedeemableNoncontrollingInterests	−0.0101	Negative
122	grossProfit	−0.0108	Negative
123	nonCurrentAssetsTotal	−0.0110	Negative
124	FCF_to_Assets	−0.0110	Negative
125	depreciation	−0.0123	Negative
126	salePurchaseOfStock	−0.0123	Negative
127	Operating_Margin	−0.0128	Negative
128	totalRevenue	−0.0129	Negative
129	netBorrowings	−0.0130	Negative
130	shortLongTermDebtTotal	−0.0132	Negative
131	ebit	−0.0137	Negative
132	totalPermanentEquity	−0.0140	Negative
133	incomeTaxExpense	−0.0145	Negative
134	sellingAndMarketingExpenses	−0.0152	Negative
135	year	−0.0153	Negative
136	taxProvision	−0.0164	Negative

Table A5. *Cont.*

Rank	Feature	Weight	Signal
137	netDebt	−0.0169	Negative
138	ebitda_TTM	−0.0170	Negative
139	netIncome_TTM	−0.0170	Negative
140	NetDebt_to_Assets	−0.0177	Negative
141	Debt_to_Assets	−0.0193	Negative
142	netInvestedCapital	−0.0193	Negative
143	noncontrollingInterestInConsolidatedEntity	−0.0205	Negative

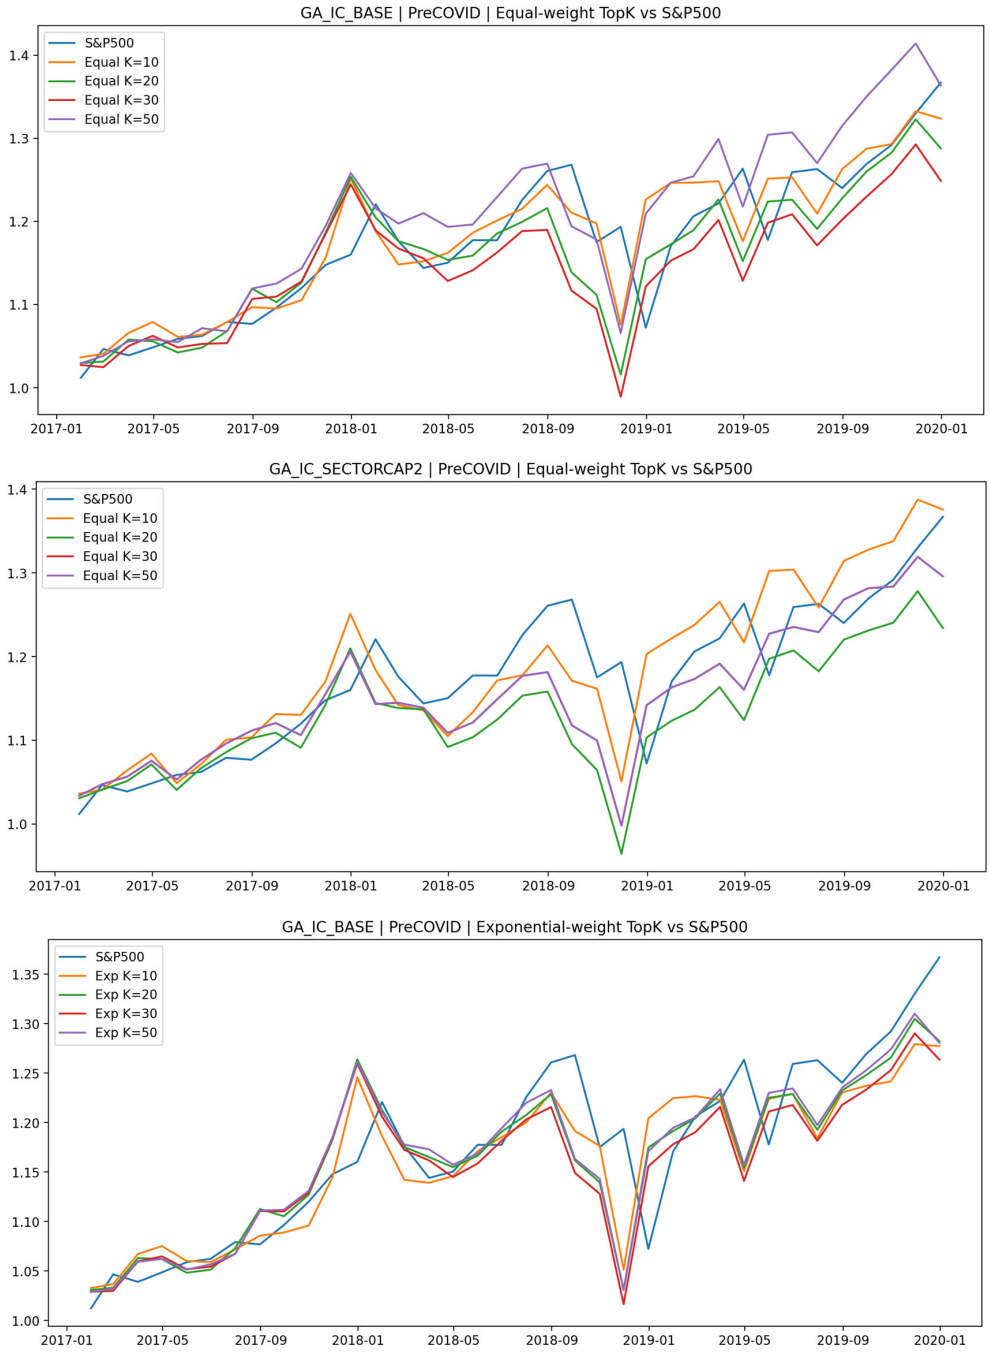


Figure A1. *Cont.*

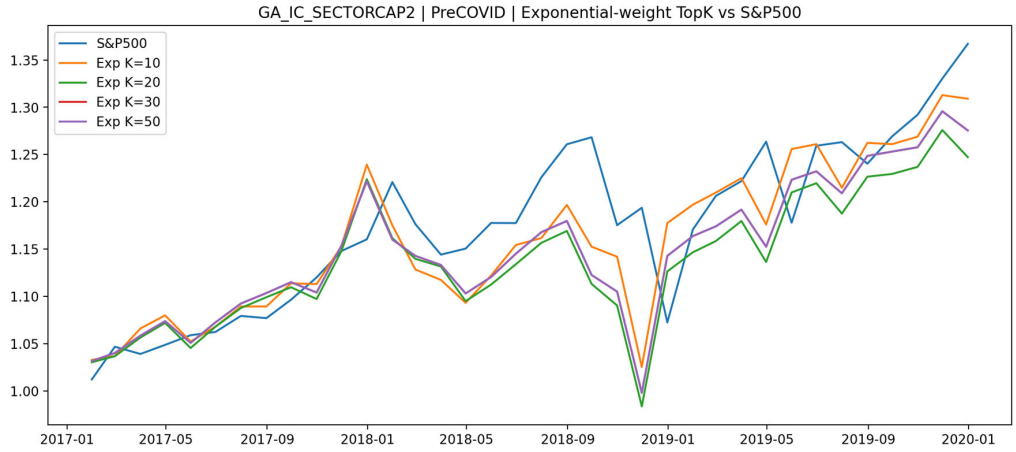


Figure A1. Pre-Covid regime equity curves—Equal and Exponational weight portfolios for GA IC.

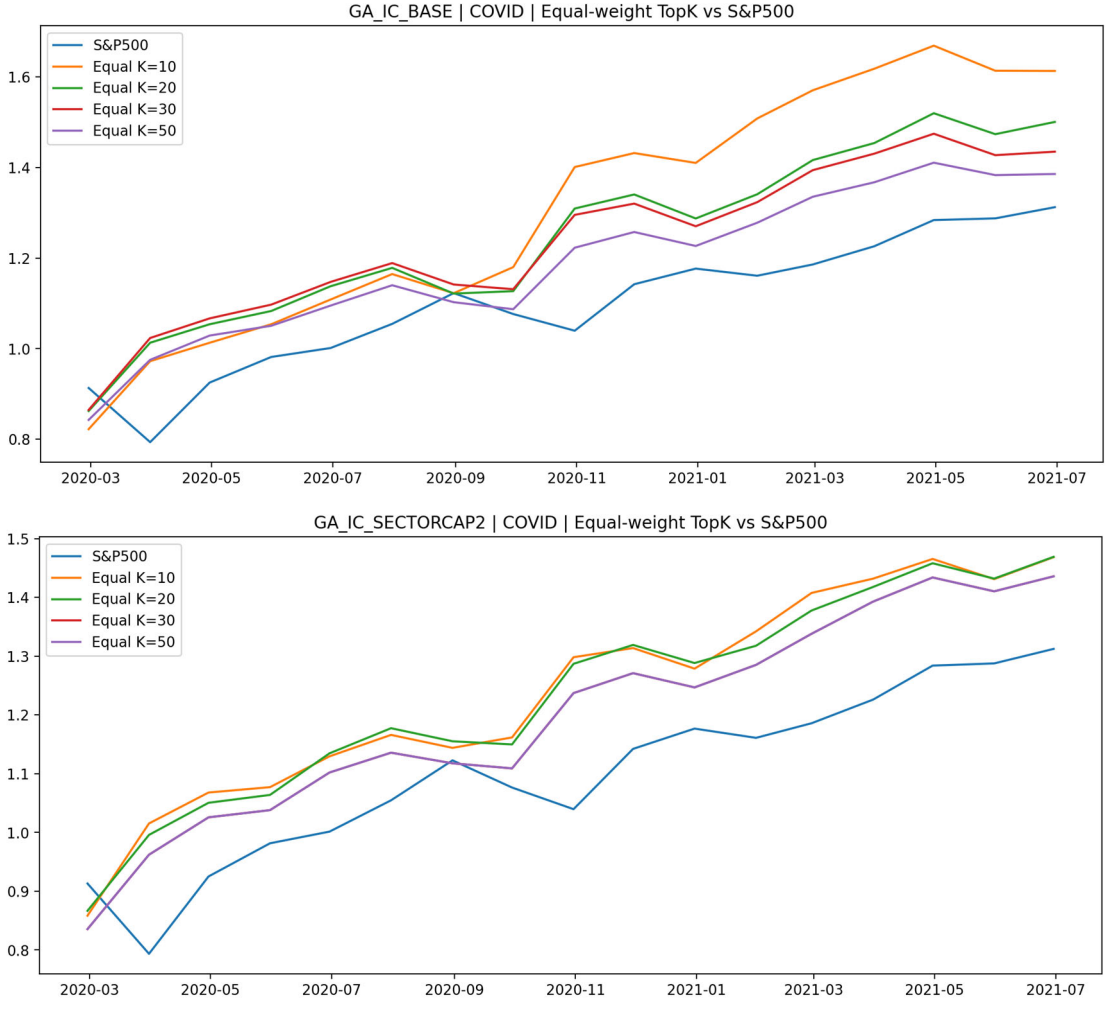


Figure A2. *Cont.*

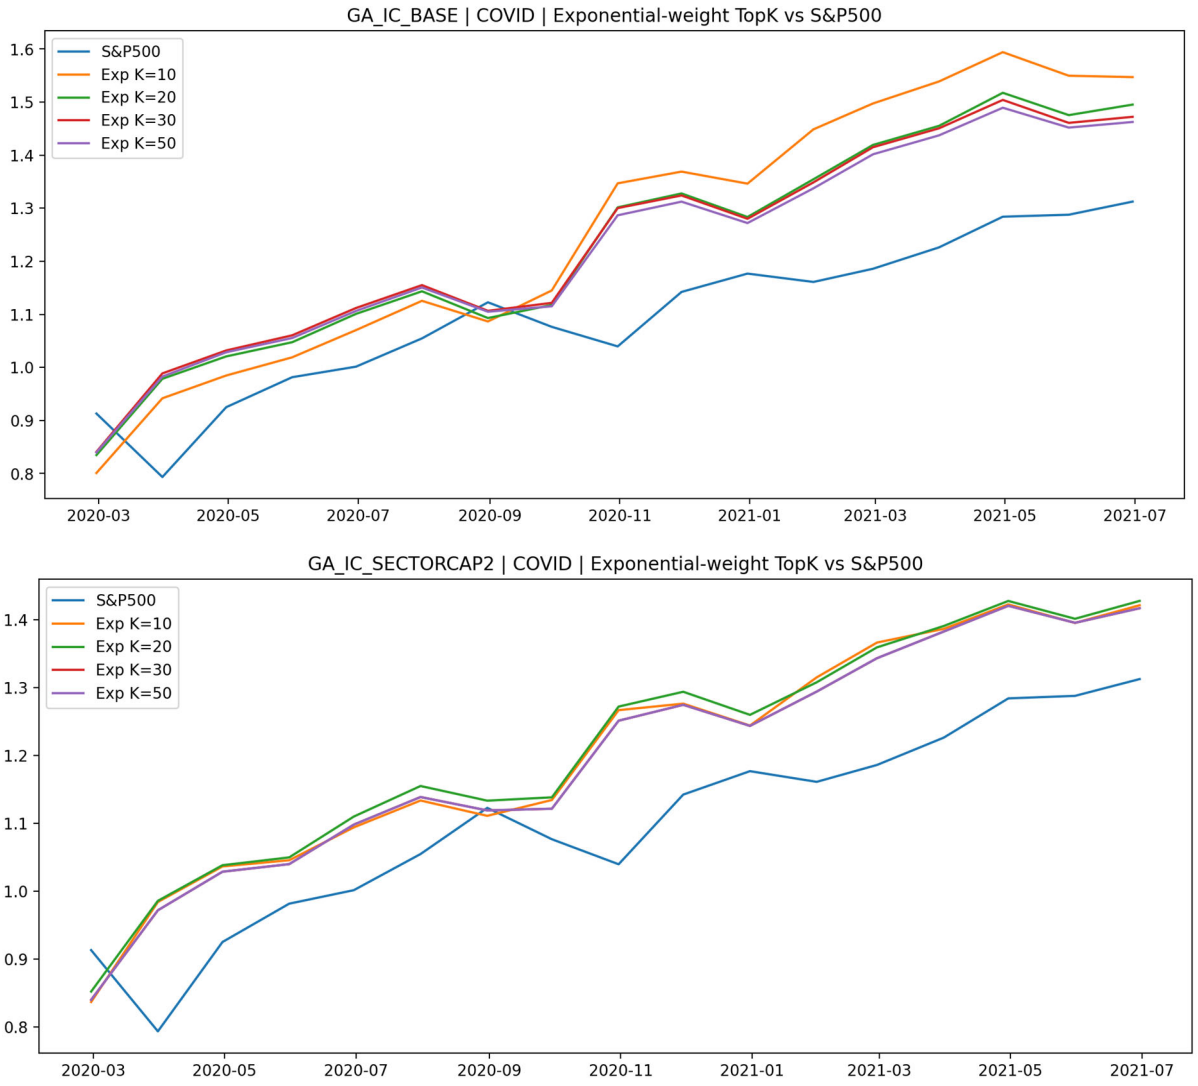


Figure A2. Covid regime equity curves—Equal and Exponational weight portfolios.

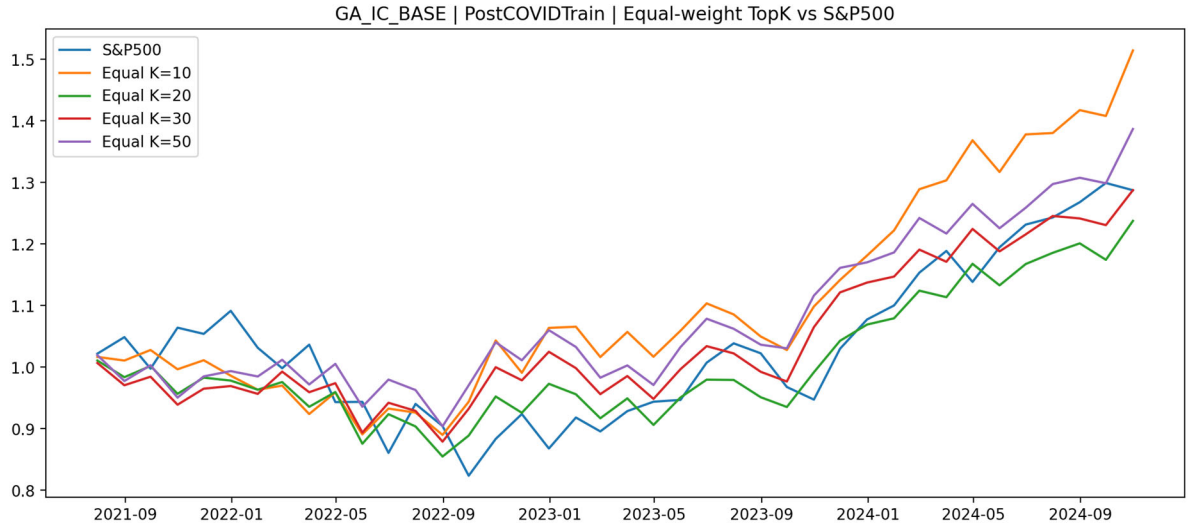


Figure A3. *Cont.*

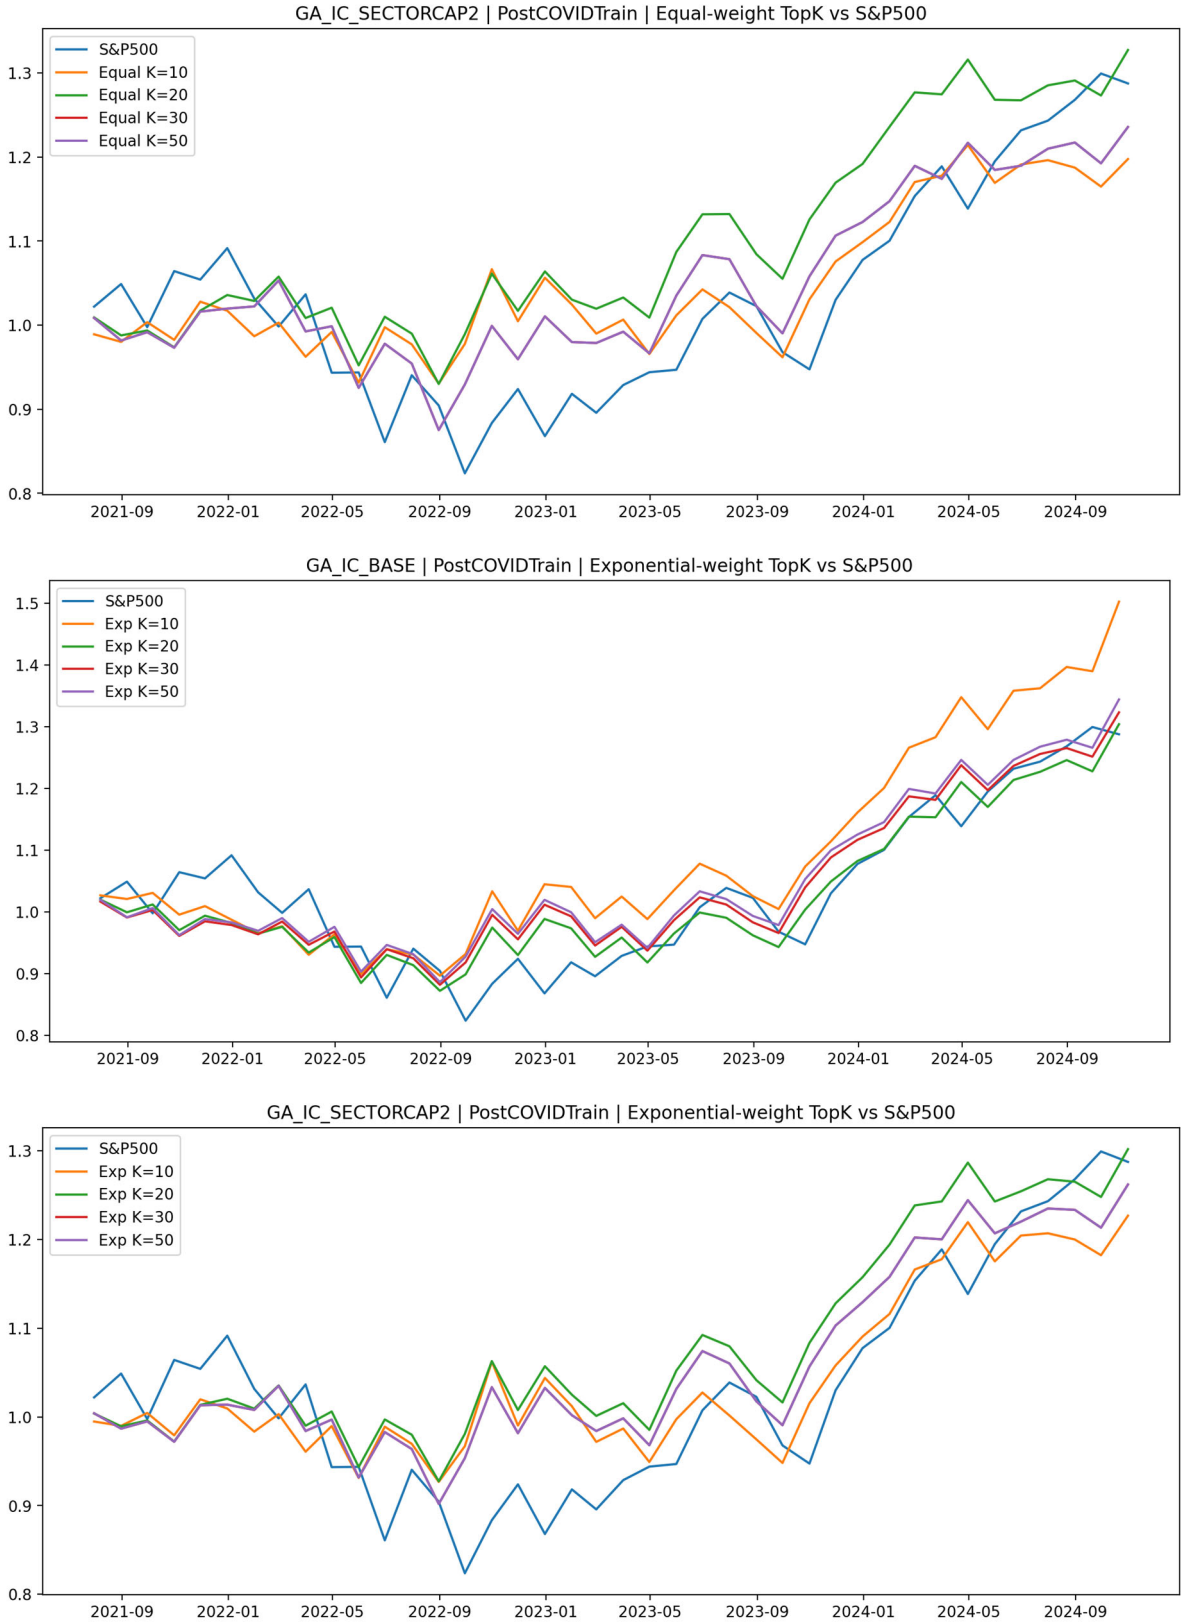
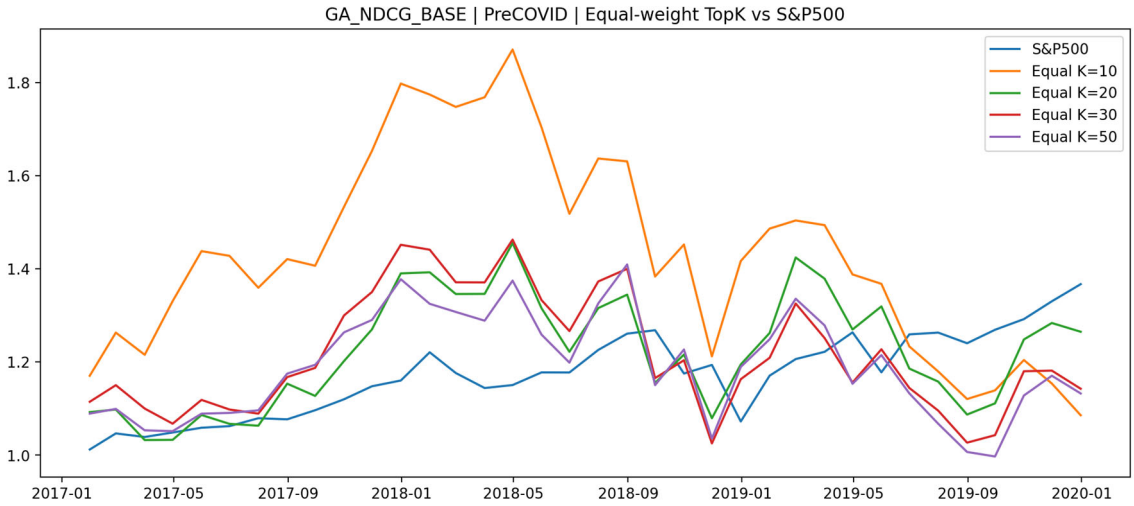
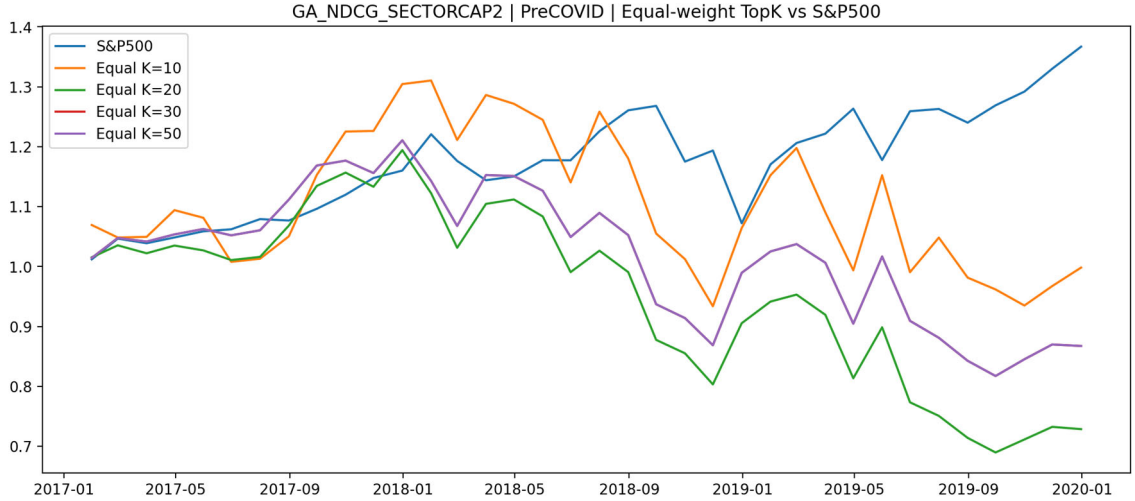


Figure A3. Post-Covid regime equity curves—Equal and Exponational weight portfolios.





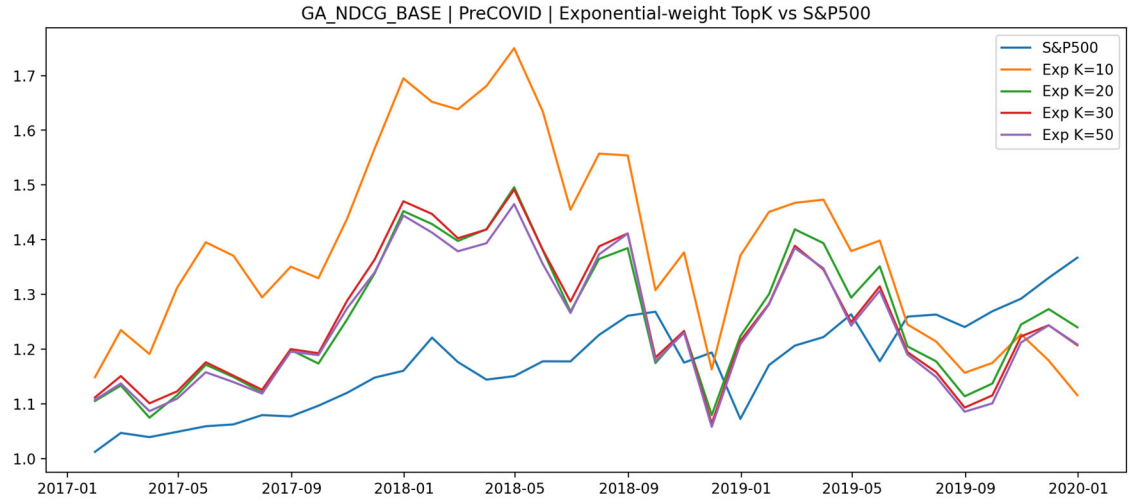


Figure A4. *Cont.*

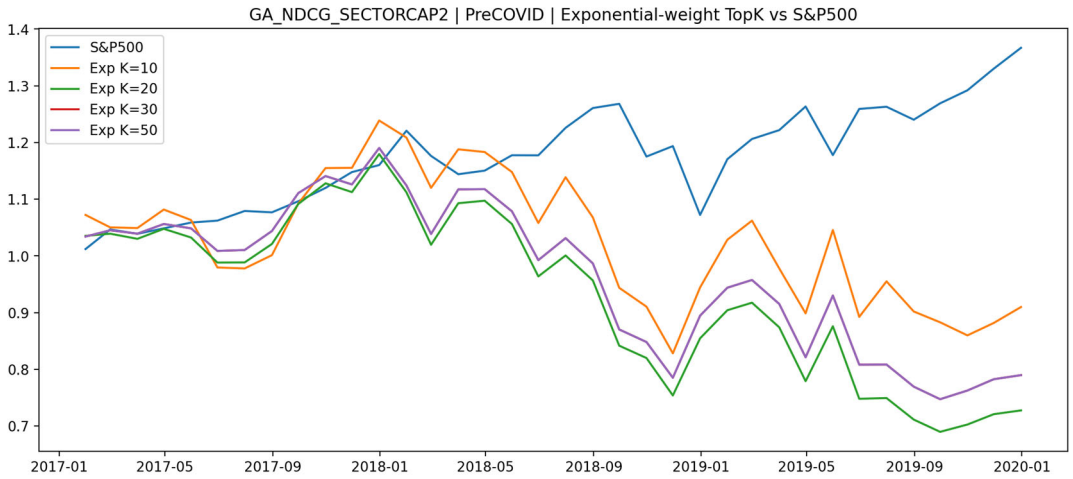


Figure A4. Pre-Covid regime equity curves—Equal and Exponational weight portfolios for GA NDCG.



Figure A5. *Cont.*



Figure A5. Covid regime equity curves—Equal and Exponational weight portfolios for GA NDCG.

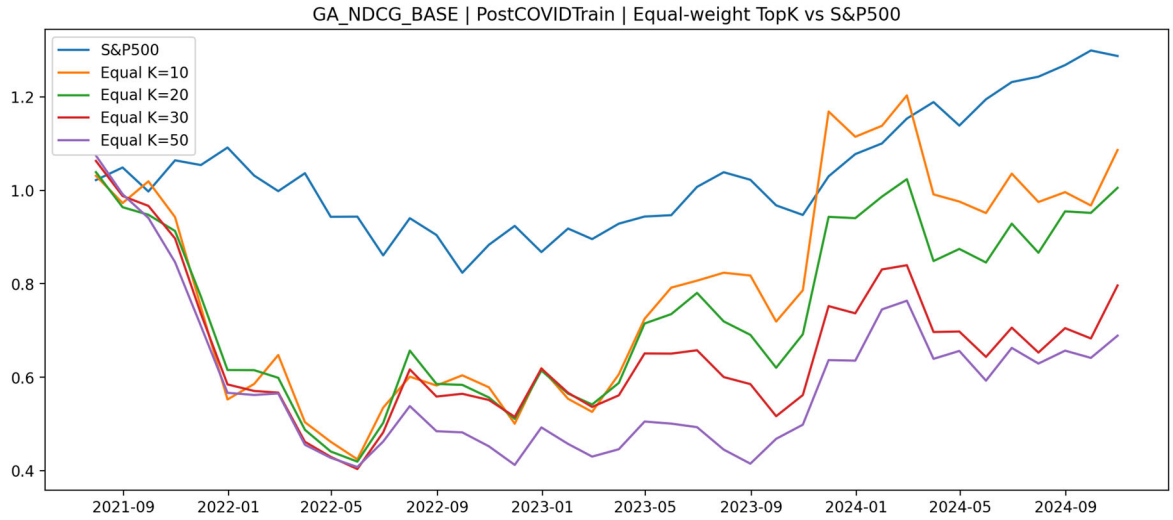


Figure A6. *Cont.*



Figure A6. Post-Covid regime equity curves—Equal and Exponational weight portfolios for GA NDCG.

References

Agarwalla, A. (2024). Enhancing stock return predictions: Comparing machine learning methods with traditional financial models. *International Journal of Social Science and Economic Research*, *9*(11), 5215–5228. [CrossRef]

Baker, S. R., Bloom, N., Davis, S. J., Kost, K., Sammon, M., & Viratyosin, T. (2020). The unprecedented stock market reaction to COVID-19. *The Review of Asset Pricing Studies*, *10*(4), 742–758. [CrossRef]

Bernard, V. L., & Thomas, J. K. (1990). Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics*, *13*(4), 305–340. [CrossRef]

Capponi, A., Huang, C., Sidaoui, J. A., Wang, K., & Zou, J. (2025). The nonstationarity-complexity tradeoff in return prediction. *arXiv*, arXiv:2512.23596. [CrossRef]

Carrasco, M., & Kon, N. (2023). Test for trading costs effect in a portfolio selection problem with recursive utility. *Journal of Financial Econometrics*, *22*(4), 908–953. [CrossRef]

Cooper, M. J., Gulen, H., & Schill, M. J. (2008). Asset growth and the cross-section of stock returns. *The Journal of Finance*, *63*(4), 1609–1651. [CrossRef]

DeMiguel, V., Garlappi, L., & Uppal, R. (2009). Optimal versus naive diversification. *Review of Financial Studies*, *22*(5), 1915–1953. [CrossRef]

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *The Journal of Finance*, *47*(2), 427–465. [CrossRef]

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, *116*(1), 1–22. [CrossRef]

Flores-Fernandez, G. A., Jimenez-Carrion, M., Gutierrez, F., & Sanchez-Ancajima, R. A. (2024). Genetic algorithm and LSTM artificial neural network for investment portfolio optimization. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, *15*(2), 27–46. [CrossRef]

Fransisca, D. C., Sukono, Chaerani, D., & Halim, N. (2024). Robust portfolio mean-variance optimization for capital allocation in stock investment using the genetic algorithm: A systematic literature review. *Computation*, *12*(8), 166. [CrossRef]

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley.

Green, J., Hand, J. R. M., & Zhang, F. (2017). The characteristics that provide independent information about average U.S. monthly stock returns. *Review of Financial Studies*, *30*(12), 4389–4436. [CrossRef]

Grinold, R. C., & Kahn, R. N. (2000). *Active portfolio management* (2nd ed.). McGraw-Hill.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, *33*(5), 2223–2273. [CrossRef]

Gunjan, A., & Bhattacharyya, S. (2024). Quantum-inspired meta-heuristic approaches for a constrained portfolio optimization problem. *Evolutionary Intelligence*, *17*(4), 3061–3100. [CrossRef]

Hou, K., Xue, C., & Zhang, L. (2015). Digesting anomalies: An investment approach. *The Review of Financial Studies*, *28*(3), 650–705. [CrossRef]

Järvelin, K., & Kekäläinen, J. (2002). Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems*, *20*(4), 422–446. [CrossRef]

Kwiatkowski, J., & Chudziak, J. A. (2025, November 10–14). *On evaluating loss functions for stock ranking: An empirical analysis with transformer model*. International Conference on Information and Knowledge Management, Seoul, Republic of Korea.

Li, J. (2023). Transaction-cost-aware random-forest portfolio construction with cardinality and minimum-weight constraints. *Bulletin of Computer and Data Sciences*, *4*(3), 1–9. [CrossRef]

Liu, T. Y. (2009). Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, *3*(3), 225–331. [CrossRef]

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

Mondal, S., Mohanty, S., Harlander, B., Kseoglu, M., Rane, L., Romanov, K., Liu, W.-K., Hatwar, P., Salath, M., & Byrum, J. (2019). Investment ranking challenge: Identifying the best performing stocks based on their semi-annual returns. *arXiv*, arXiv:1906.08636. [CrossRef]

Moskowitz, T. J., & Grinblatt, M. (1999). Do industries explain momentum? *The Journal of Finance*, *54*(4), 1249–1290. [CrossRef]

Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, *108*(1), 1–28. [CrossRef]

Nufus, S. A., Sutarman, S., & Herawati, E. (2025). Hybrid genetic algorithm for dynamic portfolio optimization problems. *SinkrOn*, *9*(3), 1777–1787. [CrossRef]

Penman, S. H. (2013). *Financial statement analysis and security valuation* (5th ed.). McGraw-Hill.

Pobrotyn, P., & Bialobrzeski, R. (2021). NeuralNDCG: Direct optimisation of a ranking metric via differentiable relaxation of sorting. *arXiv*, arXiv:2102.07831. [CrossRef]

Poh, D., Lim, B., Zohren, S., & Roberts, S. J. (2021). Building cross-sectional systematic strategies by learning to rank. *The Journal of Financial Data Science*, *3*(2), 70–86. [CrossRef]

Qiao, W., Bu, D., Gibberd, A., Liao, Y., Wen, T., & Li, E. (2023). When time varying volatility meets transaction cost in portfolio selection. *Journal of Empirical Finance*, *73*, 220–237. [CrossRef]

Sakagawa, S., & Mori, N. (2022, December 12–14). *Neural ranking strategy for portfolio construction using transformers*. 2022 13th International Congress on Advanced Applied Informatics Winter (IIAI-AAI-Winter), Phuket, Thailand.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.