Designing Agentic AI-Based Screening for Portfolio Investment

Mehmet Caner* Agostino Capponi† Nathan Sun‡ Jonathan Y. Tan§

March 25, 2026

Abstract

We introduce a new agentic artificial intelligence (AI) platform for portfolio management. Our architecture consists of three layers. First, two large language model (LLM) agents are assigned specialized tasks: one agent screens for firms with desirable fundamentals, while a sentiment analysis agent screens for firms with desirable news. Second, these agents deliberate to generate and agree upon buy and sell signals from a large portfolio, substantially narrowing the pool of candidate assets. Finally, we apply a high-dimensional precision matrix estimation procedure to determine optimal portfolio weights. A defining theoretical feature of our framework is that the number of assets in the portfolio is itself a random variable, realized through the screening process. We introduce the concept of *sensible screening* and establish that, under mild screening errors, the squared Sharpe ratio of the screened portfolio consistently estimates its target. Empirically, our method achieves superior Sharpe ratios relative to an unscreened baseline portfolio and to conventional screening approaches, evaluated on S&P 500 data over the period 2020–2024.

Keywords: Portfolio screening, Agentic design, Sharpe-Ratio, precision matrix estimation.

*North Carolina State University, Nelson Hall, Department of Economics, NC 27695. Email: mcaner@ncsu.edu.

†Columbia University. Department of Industrial Engineering and Operations Research and Columbia Business School. Email: ac3827@columbia.edu.

‡Columbia University. Department of Industrial Engineering and Operations Research. Email: nathan.sun@columbia.edu.

§Columbia University. Department of Industrial Engineering and Operations Research. Email: jyt2123@columbia.edu.

> *"We need new AI models for the real world—quantitative models... AI for the quantitative world is something else entirely, focusing on creating novel medical treatments, de novo material science, and advanced risk management and* ***portfolio construction****."*
>
> — Jack D. Hidary, CEO of Sandbox AQ. *The Wall Street Journal*, Opinion: "America Needs AI that can do Math," February 17, 2026.

> *"Ilya Sutskever has recently been adding his voice to the beyond LLM camp, stating that* ***we won't reach artificial general intelligence with large language models alone****."*
>
> — Jack D. Hidary, CEO of Sandbox AQ. *The Wall Street Journal*, Opinion: "America Needs AI that can do Math," February 17, 2026.

1 Introduction

Portfolio formation is a foundational concept in the finance literature. There are several well-known portfolio formations, including the global minimum variance portfolio, the Markowitz mean-variance portfolio, and the maximum Sharpe ratio portfolio. Traditionally, these frameworks were developed and implemented by human analysts throughout the twentieth century. With the advent of large language models coupled with deep learning-based estimation in the twenty-first century, there are now artificial intelligence-driven portfolios.

Very recently, multi-agent AI systems have been applied to portfolio optimization. This paper contributes to this emerging literature by developing a high-dimensional portfolio management framework that integrates stock screening with portfolio optimization. Every investment decision involves two distinct tasks: selecting which stocks to hold and determining their optimal weights. Accordingly, both the set of selected stocks and their portfolio weights are treated as endogenous outputs of our framework. Rather than applying quantitative methods to the full stock universe or assigning equal weights to an arbitrary subset, we propose an initial screening stage in which LLM-based agents generate buy and sell signals. A quantitative precision matrix estimation method then determines optimal investment weights for the screened stocks. This two-stage approach yields portfolios concentrated in stocks with strong return-generating potential, achieving superior Sharpe ratios relative to standard benchmarks. By incorporating stock screening, the framework deliberately departs from index mimicking, aiming instead to generate excess returns, analogous to adopting a large tracking error relative to a benchmark index in order to enhance risk-adjusted performance.

A useful analogy is the selection process of a sports team. After observing practice sessions, the manager chooses the best players to form the starting lineup for game day. The choice of the optimal subset of players can dramatically improve team performance, and similarly, selecting the optimal subset of stocks may achieve higher performance than using the entire universe. Note that

by setting up tracking error and weight constraints on portfolios, both industry practitioners and regulators move away from using all universe-indexing strategies. A recent paper by Caner and Fan (2026) outlines how to form portfolios based on data-dependent inequality constraints and establishes high-dimensional consistency of these portfolios. They demonstrate that the Sharpe ratio, returns, and collective risk of a subset of stocks can do better than those of a benchmark index. Because their restrictions are given by mutual fund prospectuses and government regulators, there is no performance-based screening involved in forming their portfolios. A study by Arvanitis et al. (2024) presents strong economic and financial theoretical foundations explaining why a sparse portfolio can perform better than or equal to the market index.

Our multi-agent AI framework consists of three coordinated layers. In the first layer, two specialized LLM agents operate in parallel: an LLM-Strategy agent (LLM-S), which screens stocks based on fundamental firm characteristics, including log firm size, book-to-market ratio, and twelve-month momentum, and a FinBERT-based sentiment agent, which analyzes financial news articles to generate monthly sentiment-driven signals. These two agents operate on complementary frequencies: LLM-S is retrained annually, capturing slow-moving structural narratives about firm quality, while FinBERT is updated monthly to respond to fast-moving news sentiment. In the second layer, the two agents deliberate and reach a consensus via an intersection-based decision rule, narrowing the candidate pool from the full S&P 500 to a high-conviction subset averaging approximately 22 stocks. In the third layer, a quantitative optimization algorithm applies state-of-the-art high-dimensional precision matrix estimation techniques, such as including nodewise regression (Meinshausen and Bühlmann (2006)), residual nodewise regression (Caner et al. (2023)), POET (Fan et al. (2013)), deep learning-based methods, and nonlinear shrinkage, to determine optimal portfolio weights under global minimum variance, mean-variance, and maximum Sharpe ratio objectives. The winning method-objective combination is selected based on out-of-sample Sharpe ratio performance.

We show that not every multi-agent AI configuration creates value, and that the design of the agent team is critical. We demonstrate this systematically by comparing our full multi-agent system against a rich set of alternatives: purely quantitative strategies without any screening, single-agent LLM systems (LLM-S alone or FinBERT alone), conventional screening methods (logistic regression, human analyst recommendations, and the Novy-Marx (2013) profitability-and-value screen), and hybrid systems that incorporate human analyst judgment alongside AI agents. In each case, our full Agentic AI architecture, specifically two specialized LLM screening agents plus a quantitative weighting method, dominates. We show that incorporating human analyst recommendations into the AI ensemble consistently *degrades* performance, a finding we attribute to the behavioral and emotional biases that human analysts inevitably carry into their recommendations. This result provides a concrete, quantitative answer to the practical criticism that the financial value of AI-based systems is difficult to measure.[1]

[1] *"Early adopters who rushed into AI pilots and even deployment last year often hit a wall and learned some hard lessons... Critically, they found it was* ***hard to measure financial returns.****"* — Isabelle Bousquette, *The Wall Street Journal*, Business Section, page B4: "AI software proves to be tougher sell than before," February 19, 2026.

Beyond the empirical results, we make a novel theoretical contribution to the high-dimensional portfolio literature. A defining feature of our framework is that the number of assets in the portfolio is not fixed in advance, but is itself a random variable realized through the screening process. The existing high-dimensional portfolio literature, including Fan et al. (2011), Caner et al. (2023), and Caner and Fan (2026), treats portfolio dimension as a deterministic, known sequence. We depart from this convention by introducing the concept of *sensible screening*: a screening process that, even when it makes mild errors in selecting the exact composition of the optimal portfolio, always includes the optimal stocks when selecting enough assets and never selects purely suboptimal ones. Under sensible screening and standard precision matrix consistency conditions, we establish in Theorem A.1 that the squared Sharpe ratio of the screened portfolio consistently estimates its target, even when the screening process is imperfect.

Our empirical analysis covers the S&P 500 from January 2020 to April 2024, a period encompassing noticeable events including the COVID-19 shock, the 2022 drawdown, and the strong recovery of 2023. The S&P 500 index achieves an annualized Sharpe ratio of 0.6324 over this window. Under a purely quantitative strategy without screening, only one method-objective combination (nonlinear shrinkage paired with the maximum Sharpe ratio objective) manages to beat this benchmark. By contrast, our multi-agent AI framework beats the market in all but one of the fifteen method-objective combinations we consider, with a peak annualized Sharpe ratio of 1.1867, an 88% improvement over the index. We extend the evaluation to a nearly ten-year horizon spanning January 2015 to April 2024, and find that the Agentic AI architecture continues to dominate, achieving a peak Sharpe ratio of 0.9429, above the market benchmark of 0.7298. These gains are accompanied by substantially higher annualized returns: the best Agentic AI configuration delivers a 36.34% annualized return over the five-year window, compared to 19.99% for the best purely quantitative baseline.

For our quantitative strategy, we follow the established literature and form weights according to global minimum variance (GMV), Markowitz mean-variance (MV), and maximum Sharpe ratio (MSR) objectives. Our quantitative estimation method applies a matrix-based investing strategy, where estimation techniques form the rows and different portfolio objectives form the columns. The winner within this matrix is the technique-objective combination that achieves the highest Sharpe ratio for the out-of-sample time period being evaluated. For example, in Table 1, the winner from January 2020 to April 2024 is the nonlinear shrinkage method of Ledoit and Wolf (2020) combined with the maximum Sharpe ratio portfolio.

The rest of the paper proceeds as follows. Section 2 provides background on large language model and discusses our proposed Agentic AI framework. Section 3 introduces the stock screening framework, including the LLM-S and FinBERT agents, alternative screening benchmarks, the algorithm, and our theoretical result on screened portfolio. Section 4 describes the quantitative high-dimensional portfolio weight formation methods. Section 5 presents our empirical results, including the main five-year evaluation, robustness checks, and the extended ten-year analysis. We

conclude in Section 6. Appendix A contains technical proofs, Appendix B provides details on the precision matrix estimation techniques, and Appendix C presents additional material including extended results, long-short strategies, LLM-S prompts and outputs, and Novy-Marx screening comparisons.

2 Agentic AI Portfolio Management Framework

We provide background on large language models and their applications in finance. We then describe the architecture of our proposed multi-agent portfolio management system and explain how it addresses the main methodological concerns raised in the literature.

2.1 Large Language Models

Unlike their predecessors, LLMs exhibit a human-like quality of text generation and deep contextual reasoning. These capabilities stem from the transformer architecture (Vaswani et al., 2017), which leverages self-attention mechanisms to capture long-range dependencies in text far more effectively than earlier recurrent or convolutional approaches.[2]

LLMs are trained on massive datasets, optimizing a vast number of parameters using deep learning estimators within a transformer structure, leading to significant performance improvements across a range of tasks (Bishop and Bishop, 2024, pp. 390–391). Leveraging self-supervised learning, LLMs acquire high-quality representations on unlabeled data and can then be further fine-tuned to improve performance. Fine-tuning is typically achieved through Low-Rank Adaptation (LoRA) or Reinforcement Learning from Human Feedback (RLHF). Alternative approaches include training the LLM from scratch on in-sample data to predict out-of-sample outcomes, or employing instruction fine-tuning, where the model is trained on specifically generated prompt-response pairs. Fine-tuned LLMs can substantially outperform generic open-source LLMs in specialized domains such as sentiment analysis and financial classification (Li et al., 2024, Section 4.1).

An important practical consideration is the choice between open-source models, such as LLaMA, and closed-source models, such as GPT-4 or Claude. While open-source models offer advantages in transparency, reproducibility, and data control, Li et al. (2024) (Section 3) find that closed-source models currently achieve superior performance on standard financial benchmarks, a gap that likely reflects differences in both model scale and the proprietary nature of their fine-tuning procedures.

[2] As documented in Kamath et al. (2024), models with fewer than ten billion parameters perform near chance level across a range of benchmarks, including arithmetic reasoning, multilingual question answering, the 57-task Massive Multitask Language Understanding (MMLU) benchmark (Hendrycks et al., 2020), and semantic understanding evaluations. Beyond this scale threshold, performance rises sharply, reaching approximately 30% on arithmetic tasks, 40–50% on multilingual question answering, and 60–70% on both MMLU and semantic understanding benchmarks.

2.2 LLMs in Finance: Related Work

LLMs offer several advantages for stock selection. Because their pre-trained nature enables robust zero-shot analysis, they eliminate the strict requirement for supervised learning. Rather than requiring costly and time-consuming model retraining, LLMs can rapidly execute simultaneous tasks such as sentiment analysis and keyword extraction, and their ability to decompose complex tasks into simpler sub-tasks makes them well suited to processing the large volumes of financial reports involved in stock selection (Li et al., 2024). A prominent empirical demonstration of this potential is provided by Chen et al. (2023), who show that integrating LLM-driven news analysis into long-short portfolios can deliver very large Sharpe ratios over the 2004–2019 period, outperforming both the S&P 500 index and classical bag-of-words approaches.

The application of multi-agent systems to portfolio construction has been considered by Zhao et al. (2025), who design a framework in which specialized LLM agents collaboratively execute portfolio construction tasks. Their system comprises three agents, a fundamental equity agent, a sentiment agent, and a valuation agent, which must collaboratively discuss, evaluate, and finalize decisions, thereby actively suppressing hallucinations made by any single agent. This multi-agent design mirrors the role of coordinated human analyst teams, which must synthesize financial disclosures, earnings calls, financial ratio analysis, market news, and research reports. Recent evidence by Zhang et al. (2025) further supports this approach: context-based LLMs outperformed their baseline counterparts by 8.6 percentage points on financial reasoning tasks (Table 2 therein), suggesting that domain adaptation substantially improves LLM performance.

Despite this promise, Li et al. (2025) raise important methodological concerns. Their central argument is that existing studies demonstrating superior LLM performance often rely on very short evaluation horizons (6–24 months) and narrow asset universes (5 to 100 stocks; see their Table 1). When evaluated over extended horizons and across a larger cross-section of stocks, they find that LLM advantages disappear, allowing simple benchmarks such as buy-and-hold to dominate. They identify three specific sources of bias. First, *look-ahead bias*, whereby the model inadvertently processes future data during training or testing. Second, *data snooping*, arising from repeated testing on a static dataset leading to inflated results. Third, *survivorship bias*, whereby delisted stocks are omitted, inflating apparent returns. Evidence of look-ahead bias is also documented by Ludwig et al. (2025) and, in the specific context of stock market predictions using annual US data over 2020–2023, by Didisheim et al. (2025), who find a small but detectable look-ahead bias at annual frequency and a substantially larger bias at high frequency.

2.3 Our Multi-Agent Portfolio Management Framework

Our framework is a two-stage pipeline applied to the full S&P 500 universe. In the first stage, an AI screening team, consisting of LLM-S and FinBERT operating in concert, identifies a targeted subset of stocks by combining fundamental analysis with sentiment analysis of news articles. In the second stage, a quantitative precision matrix estimation algorithm assigns optimal portfolio

weights to the screened stocks. This separation of screening and weighting represents a scalable, end-to-end application of Agentic AI to portfolio construction, and is the primary architectural distinction from prior work such as Zhao et al. (2025), whose study evaluates a randomly selected basket of 15 technology stocks rather than a comprehensive market-wide pipeline.

We take the following steps to prevent look-ahead bias. First, the dataset provided to the LLM-S agent contains no returns data, preventing the agent from explicitly optimizing for future performance. Second, for each retraining date, the agent queries only firm characteristics strictly available on that date, so future characteristics cannot enter the rule creation phase. Third, after the agent creates an investment rule, we manually apply it to S&P 500 firms to generate buy and sell recommendations, preventing any dynamic adjustment the agent might otherwise make. Fourth, we explicitly instruct the agent to use causal masking, which prevents the use of future test data during training. Fifth, as a robustness check, we verify that buy signals in each month do not systematically align with that month's subsequent returns. We have not observed such alignment in our analysis. Note also that our FinBERT model's weights were fine-tuned on the Financial PhraseBank dataset (Malo et al., 2014), published in 2014, giving it a knowledge cutoff date prior to 2015.

To suppress hallucinations, we deploy a consensus-based decision rule across agents. For each stock, it is included in the final recommendation only if both agents independently agree to buy or sell it; disagreement leads to exclusion. If the agents' decisions are mutually exclusive across all stocks—preventing a null recommendation—we default to the union of their recommendations. With three agents, the rule becomes a two-out-of-three majority vote. A detailed description of the decision rule is provided in Section 5.2. This intersection strategy reduces the hallucination risk inherent in any single agent while preserving the complementarity between fundamental and sentiment signals. Sapkota et al. (2026) provide a systematic comparison of single-agent LLMs and multi-agent systems, documenting that the defining advantage of multi-agent architectures lies in dynamic task decomposition, whereby a high-level objective is broken down and distributed among specialized agents. Our framework embodies this principle: LLM-S and FinBERT each address a distinct analytical sub-problem, and their outputs are combined through a formal consensus mechanism rather than informal aggregation.

Following Li et al. (2025), we prevent data snooping by using rolling windows rather than a fixed in-sample training period. We address survivorship bias by including all stocks in the S&P 500, including those subsequently delisted.

To address the short-horizon critique of Li et al. (2025), we analyze two evaluation periods. The first spans five years and encompasses a diverse set of market regimes: the COVID-19 shock, the 2022 market drawdown, and the strong recovery of 2023. The second extends the analysis to a nearly ten-year horizon from 2015 to 2024.

3 Stock Screening and Weights

Having outlined the architecture of our multi-agent system, we now develop its core component in detail. We formalize the concept of stock screening, describe how our LLM agents implement it, present alternative screening benchmarks, and establish a theoretical result on the Sharpe ratio consistency of screened high-dimensional portfolios

One argument for stock screening is the identification of potentially high-return assets within a large universe of stocks. Additionally, stock screening can help prevent behavioral biases which are common among human analysts. The value of screening has been highlighted in several studies, starting with the foundational work of Fama and French (1992), who demonstrate that screening for small-cap and low price-to-book stocks leads to higher returns in the twentieth century. Jegadeesh and Titman (1993) validate the idea of "price-performance" screening and consider the "momentum" effect, showing that stocks performing well over the past 3-12 months tend to continue performing well in the near future. Piotroski (2000) introduces F-scores: nine accounting ratios to find high quality stocks. Novy-Marx (2013) adds a quality-gross profitability criterion for value stocks.

In sum, screening generates higher returns and/or higher Sharpe Ratio portfolios for several reasons. First, due to the limited cognitive capacity of human beings, technical analysis may be more effective on a subset of stocks compared to, for example, a larger universe of a thousand stocks. Second, since investors are susceptible to narrative bias, overweighting compelling stories attached to assets rather than their fundamentals (Tversky and Kahneman, 1974, 1983), a systematic, quantitative screening approach can help identify and exclude stocks driven primarily by sentiment rather than intrinsic value.

Further, financially distressed firms can be systematically identified by analyzing debt and liquidity ratios and avoiding buying underperforming or close to bankruptcy firms by using the Altman Z score (see Altman (1968)). Eliminating or shorting distressed firms from the portfolio is another avenue to higher returns.

3.1 LLM-S and FinBERT

Our distinguishing approach compared to the rest of the literature is that we provide a multi-agent AI system that can screen stocks across a broader universe.

We start with developing our own LLM agent, tailored at analyzing firm fundamental data. We call this agent LLM-S, abbreviated for LLM-strategy. Using standard AI agent construction techniques, LLM-S produces a fundamentals-based scoring rule to screen stocks. The objective is to filter out firms that have undesirable fundamentals, leaving behind only high-quality stocks in the portfolio. We utilize LLMs in a novel way: we prompt the LLM to adopt the persona of a portfolio manager, and we ultimately let the LLM decide upon a specific buy and sell strategy based on fundamentals data. In doing so, we also gain insight into its explanation on why it has

chosen a specific buy and sell strategy. At the end of this choice, LLM-S is asked to justify its screening choice.

We next describe the main steps of our LLM-S agent. First, LLM-S uses three factors to initiate the algorithm: log firm size (mve), book-to-market ratio (bm), and 12-month momentum (mom12m). We select these specific factors because they represent the most robust and widely accepted predictors of cross-sectional stock returns in the empirical finance literature, capturing the size of the firm, its valuation, and its recent performance, respectively (see Fama and French (1992) and Carhart (1997)). Furthermore, our preliminary testing revealed that using this parsimonious set of factors is crucial for the LLM. Expanding the factor set to include the Fama-French 5-factor variables plus momentum empirically degraded the model's performance. So it is critical to model the right LLM agent rather than simple use of LLM. This is very much related to architecture of LLM agents and the aim here is to find agents that add economic or financial value to the firm. This design of the architecture is crucial to reflect criticism in the practical financial world about the value of LLM (see Footnote 1 in Introduction). The LLM-S is instructed to follow the following four steps on its own without supervision, and we explicitly instruct it to use causal masking to prevent look-ahead bias.

1. Explore the data to understand extreme values in factors, clustering or breakpoints, and correlations between factors.
2. Develop clear rules based on economic intuition.
3. Define specific thresholds for buy and sell.
4. Provide rationale for decisions.

Each of these points is extensively detailed in Appendix C, where a snippet of our code is also shown. These above four rules can be thought of the following pipeline: data analysis, developing rules for buying and selling, deciding on exact thresholds, and finally, an explanation of its decision rationale. By using size, book-to-market, and momentum as features, our LLM-S filters the universe of stocks to retain only those exhibiting characteristics historically associated with risk premia. The strategy also inherently acts as a signal enhancer: by isolating firms with the most extreme signals, LLM-S constructs a subset of firms with high factor exposure. This provides the quantitative estimation method a highly informative candidate pool to optimize upon.

Our LLM-S deviates from traditional screening methods in a few fundamental ways. Firstly, traditional methods often rely on fixed heuristic rules that are static over time (see, for example, Piotroski (2000) and Mohanram (2005)). In contrast, our LLM-S is rerun once every year to ensure the most recent rules, and it has the ability to choose exact thresholds (i.e. it has the ability to contextualize the rule based on the distribution of the data for a specific date). While previous methods mostly rely on backward-looking training to determine scoring rules (see, for example,

Altman (1968) and Asness et al. (2019)), we leverage the broad knowledge embedded in pre-trained modern LLMs. Because these models have already learned extensive financial principles from a massive set of data, we can use a "zero-shot" approach. That is, the model generates the scoring rule directly from its internal knowledge without requiring any training. Lastly, the LLM provides the scoring rule *plus* the economic intuition behind it, providing a greater level of interpretability.

As an example, our LLM-S provided the following strategy and rationale for buy and sell decisions in year 2024 (more details regarding reasoning of this choice are in Appendix C).

- BUY Rule: This rule targets undervalued (high bm), reasonably sized ($mve > 0.3$) companies with positive momentum ($mom12m > -0.5$).[3] The economic intuition is to buy companies that are currently cheap but have shown some signs of recovery or positive market sentiment.

- SELL Rule: This rule aims to sell companies that are overvalued (low bm), have negative momentum ($mom12m < -0.55$), or are small in size ($mve < -0.75$).

Note that buy and sell rules underscore the presence of interaction effects within financial markets, aligning with the theoretical framework established in Section 3.6 of Kelly and Xiu (2023). In our example above, our buy decision is reminiscent of the interaction term between high book-to-market, reasonably sized, and positive momentum firms. Our sell decision can be similarly viewed as an interaction term. In sum, our buy and sell decisions are seen as analyzing the buy and sell decisions from a nonlinear framework involving interaction between factors. Our results can also be related to an information theoretic approach by Hong et al. (2000), who demonstrates that the momentum effect interacts with size of the firm and analyst coverage.

After introducing LLM-S, to handle human biases and account for short-term news, we utilize FinBERT, a natural language processing model specialized in analyzing financial text, to conduct sentiment analysis on recent stock news. FinBERT agent also makes buy-sell decisions. After discussing their buy and sell decisions, two agents decide on a consensus of buy-sell decisions.

Once these LLMs decide which stocks to buy and sell, we feed the screened stocks into quantitative high-dimensional weight formation techniques (see Callot et al. (2021), Caner et al. (2023) using nodewise or residual nodewise regression respectively). These methods will be described in detail in subsequent sections. In short, we combine LLM- and sentiment analysis-based screening with state-of-the-art high dimensional quantitative portfolio formation techniques.

We also connect our screening idea econometrically to high-dimensional portfolio analysis. The screening process can be conceptualized not just as selecting stocks that can provide higher returns or Sharpe ratios, but it can also be viewed as choosing the correct number of stocks among a large universe of them. Clearly, this screening choice is a random variable and affects how we choose the weights in our portfolio. First, let us denote the number of stocks chosen by the screening process as the random variable $\hat{p}$. Also, denote by p^* the target (optimal) number of stocks, which is non-random, and $p \geq p^* \geq 1$, where p is the size of the entire universe of stocks. It is worth

[3] Note the mve, bm, and mom12m features are standardized to have mean 0 and variance 1.

emphasizing that, together with the number of stocks, we have the stocks themselves; for example, the optimal choice may be Nvidia and Palantir; hence giving us $p^* = 2$. We have the following definition of the screening process. Let $\mathcal{U}$ represent the total universe of stocks.

Definition: Sensible Screening. Let $\mathcal{S}^* \subseteq \mathcal{U}$ denote the optimal portfolio, with $|\mathcal{S}^*| = p^*$, and let $\hat{\mathcal{S}} \subseteq \mathcal{U}$ denote the set of stocks selected by the screening process, with $|\hat{\mathcal{S}}| = \hat{p}$. A screening process is *sensible* if:

(i) when $\hat{p} \geq p^*$, the optimal set is recovered within the screened set: $\mathcal{S}^* \subseteq \hat{\mathcal{S}}$;

(ii) when $\hat{p} < p^*$, the screened set is entirely composed of optimal stocks: $\hat{\mathcal{S}} \subset \mathcal{S}^*$.

A sensible screening process never selects a stock that does not belong to the optimal portfolio, and always includes the optimal portfolio when it selects enough stocks. For example, suppose $p^* = 2$ and the optimal portfolio is $\mathcal{S}^* = \{\text{Nvidia, Palantir}\}$. If $\hat{p} = 3$, sensible screening requires that $\hat{\mathcal{S}} = \{\text{Nvidia, Palantir, X}\}$ for some third stock $\text{X} \notin \mathcal{S}^*$, i.e., the two optimal stocks are always included. If instead $\hat{p} = 1$, sensible screening requires that the single selected stock be either Nvidia or Palantir.

The main idea of our theoretical analysis is to treat the number of stocks in the portfolio as a random variable, one that is realized through the screening process before any weights are estimated. Under sensible screening, and robust to mild screening errors, we show that the Sharpe ratio of the screened portfolio consistently estimates the target Sharpe ratio of the optimal stock universe. This result appears to be new in the high-dimensional portfolio literature, where the composition of the investable set is typically taken as given. We next turn to the screening procedures used in our empirical analysis.

3.2 Different Screening Methods

As a benchmark, we also consider alternative screening methods. We consider stock picking by human analysts, and a logistical regression-based stock picking method. In the first method, we use the buy/sell decisions of human analysts from IBES (Institutional Brokers Estimates System, accessed via WRDS) recommendations, and in the second one, we use logistic regression to generate buy/sell signals as proposed by Chen et al. (2023) . The former also represents conventional human judgment, and the latter is a simple quantitative strategy for stock picking. Our benchmarking strategy is designed to disentangle three sources of potential performance gains: stock screening per se, the use of LLMs for screening, and the benefits of a multi-agent architecture. To this end, we compare our multi-agent system against four classes of alternatives. First, we benchmark against conventional screening methods, namely momentum and value screens, to assess whether any performance improvement stems from screening itself rather than from the specific screening technology. Second, we compare against single-agent LLM screeners, such as LLM-S alone or FinBERT combined with the quantitative method, to isolate the contribution of the multi-agent

architecture. Third, we consider alternative multi-agent configurations, including hybrid systems that combine LLM-S, FinBERT, and human judgment screening, with portfolio decisions delegated to a quantitative optimization method. Finally, we compare against the quantitative-method-only baseline, which abstracts from screening entirely, to quantify the marginal value of LLM-based stock selection

3.3 Algorithm

Our algorithm screens p stocks and subsequently feeds the results to a quantitative precision matrix estimation method, who determines the portfolio weights. Our approach proceeds as follows:

1. The LLM-S agent and the FinBERT agent independently decide which stocks to buy/sell. We rerun the LLM-S agent each year, and the FinBERT agent each month. We do this to capture the slow-moving economic narrative that the LLM-S aims to capture while also preserving the fast-moving news that FinBERT captures. This mirrors classical portfolio literature (see Fama and French (1992)) and modern ML approaches (see Gu et al. (2020)) that refit computationally heavy models annually.

2. According to their decision rule (which will be discussed in Subsection 5.2), the models reach a consensus, if applicable, on a subset of $\hat{p}$ stocks to buy/sell, where $1 \leq \hat{p} \leq p$.

3. These $\hat{p}$ stocks are provided to the quantitative method, which computes optimal portfolio weights based on a strategy described in Section 4. These depend on a combination of a method and objective.

4. Repeat for all months. We use the annualized out-of-sample Sharpe ratios as the primary measure of performance.

We illustrate the algorithm in Figure 1.

Denote the covariance matrix of outcomes as Σ. As an illustration of the third step, consider the case where the portfolio objective is global minimum variance (GMV) and the precision matrix, which is the inverse of the covariance matrix of outcomes, $\Gamma := \Sigma^{-1}$ is estimated via nodewise regression. The GMV weights, computed via the nodewise-GMV method, are then given by

$$\hat{w} = \frac{\hat{\Gamma} 1_{\hat{p}}}{1'_{\hat{p}} \hat{\Gamma}_{\hat{p}} 1_{\hat{p}}},$$

where $\hat{\Gamma}$ is the $\hat{p} \times \hat{p}$ estimated precision matrix of the screened stocks and $1_{\hat{p}} \in \mathbb{R}^{\hat{p}}$ is a vector of ones.

More details are described in Section 4. Note the fact that dimension of the precision matrix ($\hat{p}$) is a random variable, determined by the screening process prior to estimation. This stands in contrast to the existing high-dimensional portfolio literature, where the number of assets is treated

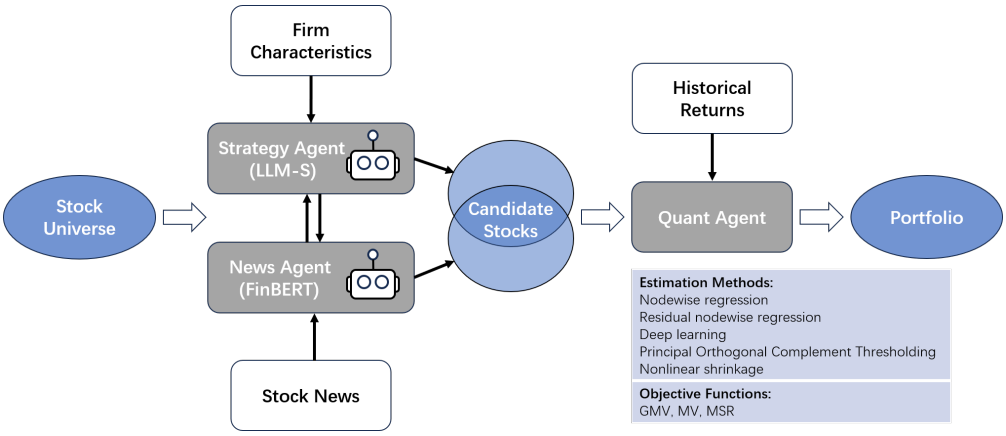


Figure 1: An illustration of the algorithm. LLM-S and FinBERT first screen for candidate stocks to invest in, and the precision matrix estimation technique assigns weights to these stocks.

as a deterministic sequence growing at a known rate; see, e.g., Fan et al. (2011), Ledoit and Wolf (2017), Caner et al. (2023), and Caner and Fan (2026).

3.4 Screened Portfolio based Sharpe Ratio Analysis

Let $\widehat{SR}_{\hat{p}}$ denote the Sharpe ratio estimator based on the screened portfolio of dimension $\hat{p}$, and let SR_{p^*} denote the target Sharpe ratio of the optimally screened portfolio. In Theorem A.1, we establish that under sensible screening with mild errors (formalized in Assumption A.1),

$$\left| \frac{\widehat{SR}^2_{\hat{p}}}{SR^2_{p^*}} - 1 \right| = o_p(1).$$

That is, the squared Sharpe ratio of the screened GMV portfolio consistently estimates its target, even when the screening process is imperfect. To our knowledge, this is the first such result in the high-dimensional portfolio literature, where portfolio dimension is typically treated as fixed and known. The result holds for any sensible screening process satisfying Assumption A.1, and is therefore compatible with a broad class of precision matrix estimators, including those of Fan et al. (2011) and Caner et al. (2023). The results will hold true in large portfolios as long as the precision matrix estimation of asset returns are consistent. Case-by-case analysis can be obtained under weaker conditions than the general level assumptions in Appendix A.

4 Quantitative Based High Dimensional Portfolio Weight Formation

One of the main empirical discoveries by Caner and Fan (2025) and Caner et al. (2023) is a joint analysis of precision matrix estimation techniques with objective functions. The main insight will

be investing in the winner of the best combination of precision matrix estimation technique and objective function. In all the quantitative-based techniques, we let the number of assets to be p, and time series to be n. We index assets by $j = 1, \cdots, p$ and index time by $t = 1, \cdots, n$.

4.1 Precision matrix estimation

In this section, we discuss methods for high dimensional portfolio selection. We include brief descriptions of five different methodologies for precision matrix estimation, representing a variety of approaches ranging from statistical shrinkage to factor models to ML-based regression techniques. These methods address the challenge of estimating the precision matrix $\Gamma = \Sigma^{-1}$ when the number of assets exceeds the sample size. We use the estimated $\hat{\Gamma}$ of these methods as the precision matrix when making portfolio decisions in future sections. In each of the methods below in Section 4, we use a constant p assets to determine the estimator for the precision matrix. In the appendix, we show how a screened universe of stocks change the precision matrix estimation. A rigorous, in-depth description of each method can be found in Appendix B.

4.1.1 Nodewise regression

Introduced by Meinshausen and Bühlmann (2006) and applied to portfolio risk-estimation by Callot et al. (2021), nodewise regression estimates the precision matrix directly via p Lasso linear regressions. By modeling each asset's excess return as a linear combination of all other assets' excess returns (i.e. $y_{t,j} = y'_{t,-j}\gamma_j + \eta_{t,j}$), this method explicitly imposes sparsity on the rows of the precision matrix. To account for high dimensions, the coefficients are estimated via a Lasso regression:

$$\hat{\gamma}_j = \text{argmin}_\gamma \left[\frac{||y_j - Y_{-j}\gamma||_2^2}{n} + 2\lambda_j ||\gamma||_1 \right],$$

where y_j is the $n \times 1$ vector of returns for asset j, Y_{-j} is the $n \times (p-1)$ matrix of returns for all other assets, and $||\cdot||_1$ and $||\cdot||_2$ denote the standard l_1 and l_2 norms. High-dimensional consistency is achieved by optimizing the penalty parameter λ_j via a Generalized Information Criterion (GIC). The precision matrix can be constructed directly by using matrix algebra and defining the diagonal elements $\hat{\Gamma}_{j,j} = \hat{\tau}_j^{-2}$ and the vector of off-diagonal elements for $j-th$ row $\hat{\Gamma}_{j,-j} = -\hat{\tau}_j^{-2}\hat{\gamma}'_j$, where $\hat{\tau}_j^2 = \frac{||y_j - Y_{-j}\hat{\gamma}_j||_2^2}{n} + \lambda_j ||\hat{\gamma}_j||_1$. We form each row j of $\hat{\Gamma}$ by using the main diagonal term $\hat{\Gamma}_{j,j}$ and the off-diagonal term in row j as $\hat{\Gamma}_{j,-j}$. Then stacking each row one upon other we form $\hat{\Gamma}$. $\hat{\Gamma}$ is then the nodewise estimator of the precision matrix.

4.1.2 Residual nodewise regression

Proposed by Caner et al. (2023), this approach extends standard nodewise regression by integrating factor models. The asset returns are modeled by $y_{t,j} = b'_j f_t + u_{t,j}$, where f_t represents a $K \times 1$ vector of observable factors (e.g., the Fama-French three-factor model) and $b_j : K \times 1$ represents

the factor loadings. Unlike standard nodewise regression, which assumes the precision matrix of returns is sparse, this method only assumes sparsity in the precision matrix of the unobserved idiosyncratic errors, Σ_u^{-1}. After the observable factor structures are removed via ordinary least squares, nodewise regression is applied to the residuals to estimate $\hat{\Omega} \approx \Sigma_u^{-1}$. The final precision matrix of returns is reconstructed analytically using the Sherman-Morrison-Woodbury formula:

$$\hat{\Gamma} = \hat{\Omega} - \hat{\Omega}\hat{B}[\hat{\Sigma}_f^{-1} + \hat{B}'\hat{\Omega}_{sym}\hat{B}]^{-1}\hat{B}'\hat{\Omega},$$

where $\hat{B}$ is the $p \times K$ matrix of estimated factor loadings, $\hat{\Sigma}_f$ is the sample covariance of the factors, and $\hat{\Omega}_{sym} = (\hat{\Omega} + \hat{\Omega}')/2$ is the symmetrized residual precision matrix estimate.

4.1.3 Principal Orthogonal Complement Thresholding (POET)

Developed by Fan et al. (2013), POET is designed for linear factor models ($y_{t,j} = b_j' f_t + u_{t,j}$) where the K common factors f_t are unobservable and must be estimated. POET first uses principal components analysis (PCA) to estimate the unobservable factors, then uses a thresholding method to estimate the covariance matrix of errors. Under the assumption that the covariance matrix of the remaining idiosyncratic errors (Σ_u) is sparse, POET applies a soft-thresholding technique to the error covariance matrix to eliminate spurious correlations, yielding $\hat{\Sigma}_{u,Th}$ (we leave technical details in Appendix B). The final precision matrix is reconstructed from the thresholded error covariance and the estimated factor structure:

$$\hat{\Gamma} = \hat{\Sigma}_{u,Th}^{-1} - \hat{\Sigma}_{u,Th}^{-1}\hat{B}(I_K + \hat{B}'\hat{\Sigma}_{u,Th}^{-1}\hat{B})^{-1}\hat{B}'\hat{\Sigma}_{u,Th}^{-1},$$

where $\hat{B}$ is the matrix of PCA-estimated factor loadings and I_K is the $K \times K$ identity matrix.

4.1.4 Deep learning

To capture complex, non-linear relationships between asset returns and observable factors, Caner and Daniele (2025) introduce a deep learning-based estimator. The asset returns are modeled by a multi-layer neural network, $y_{t,j} = g_j(f_t) + u_{t,j}$, where f_t is a K-dimensional observable column vector and $g_j(\cdot)$ is an unknown function. This effectively decomposes the total covariance matrix into a non-linear function covariance component and an idiosyncratic error covariance component ($\Sigma_y = \Sigma_g + \Sigma_u$). Similar to POET, a targeted thresholding mechanism is applied to the error covariance matrix ($\hat{\Sigma}_{u,Th}$) to ensure stability. The final precision matrix is then recovered algebraically:

$$\hat{\Gamma} = \hat{\Sigma}_{u,Th}^{-1} - \hat{\Sigma}_{u,Th}^{-1}\hat{\Sigma}_g(I_K + \hat{\Sigma}_{u,Th}^{-1}\hat{\Sigma}_g)^{-1}\hat{\Sigma}_{u,Th}^{-1},$$

where $\hat{\Sigma}_g$ is the estimated covariance matrix of the neural network predictions.

4.1.5 Nonlinear Shrinkage (NLS)

Introduced by Ledoit and Wolf (2017) and advanced by Ledoit and Wolf (2020), Nonlinear Shrinkage (NLS) addresses the instability of the sample covariance matrix in high dimensions through modifying its eigenvalues. Starting with the spectral decomposition of the sample covariance matrix $S = U\Lambda U'$ (where U is the orthogonal matrix of eigenvectors and Λ is the diagonal matrix of sample eigenvalues), NLS systematically modifies the sample eigenvalues while preserving the eigenvectors. Utilizing the Hilbert transform of the sample spectral density (elaborated in Appendix B), NLS derives an optimal, closed-form final estimator, which takes on the form:

$$\hat{\Sigma} = U\hat{\Delta}^* U',$$

where $\hat{\Delta}^*$ is the diagonal matrix of optimally shrunk eigenvalues. Conceptually, this local shrinkage pulls dispersed sample eigenvalues toward each other, correcting the systemic over-dispersion inherent in high-dimensional settings and producing a well-conditioned precision matrix estimator. We then set $\hat{\Gamma} = \hat{\Sigma}^{-1}$.

4.2 Objective functions

We introduce three objective functions that are heavily used in practice and provide the optimal weight of the portfolios attached to these functions. These are the global minimum variance portfolio (GMV), the Markowitz mean-variance portfolio (MV), and the maximum Sharpe ratio portfolio (MSR).

The GMV weights are given by

$$w^* := \text{argmin}_{w\in\mathbb{R}^p} w'\Sigma w, \quad \text{such that } w'1_p = 1.$$

The solution to the above is well-known and is given by

$$w^* = \frac{\Gamma 1_p}{1_p'\Gamma 1_p}.$$

The GMV portfolio is designed to be very risk adverse, as it is designed to minimize the variance of a portfolio. It may be a good device to provide a small return without too much risk. A different objective function that incorporates both return and variance would be the mean variance portfolio, whose weights are given by

$$w^* := \text{argmin}_{w\in\mathbb{R}^p} w'\Sigma w, \quad \text{such that } w'1_p = 1, \quad w'\mu = \rho.$$

Here, ρ is the target monthly return (in our empirics we use $\rho = 0.01$). The mean return μ is estimated by the average return in our train window. The solution is also well-known and is given

by

$$w^* = \left(\frac{D - \rho F}{AD - F^2}\right)\Gamma 1_p + \left(\frac{\rho A - F}{AD - F^2}\right)\Gamma\mu,$$

where $A := 1_p'\Gamma 1_p$, $F := 1_p'\Gamma\mu$, and $D := \mu'\Gamma\mu$.

Finally, the maximum Sharpe ratio portfolio is heavily used in practice, whose weights are given by

$$w^* := \text{argmax}_{w\in\mathbb{R}^p} \frac{w'\mu}{\sqrt{w'\Sigma w}}, \quad \text{such that } 1_p'w = 1.$$

The solution is given by

$$w^* = \frac{\Gamma\mu}{1_p'\Gamma\mu}.$$

5 Empirical Results

This section addresses two empirical questions. First, does combining LLM-based screening with a quantitative weighting method produce superior out-of-sample Sharpe ratios relative to the quantitative method operating without any screening? This isolates the incremental value of the multi-agent screening stage. Second, do conventional screening approaches, specifically logistic regression-based screening and human analyst recommendations, achieve comparable performance when paired with the same quantitative method? We show that the multi-agent AI framework outperforms both alternatives across all specifications.

We also conduct a series of robustness checks. First, we evaluate a hybrid agent system (specifically, FinBERT, LLM-S, plus human analyst recommendations) with the quantitative method to determine if incorporating human judgment enhances or detracts from the Agentic AI system. Second, we assess whether a single-agent screening approach (pure FinBERT or LLM-S) and the quantitative method will outperform the two LLM agents together plus the quantitative method. Lastly, we evaluate how human-only screening with the quantitative method performs compared with only the quantitative method. This provides a measurable assessment of the value of traditional human screening.

We report the following performance measures: the mean (out-of-sample) monthly return, the out-of-sample variance, and the out-of-sample Sharpe ratio. All of these will be calculated with a transaction cost of 10 basis points. Let $y_{P,t+1} = \hat{w}_t' y_{t+1}$ be the gross return of the portfolio at time $t+1$, and $\hat{w}_t$ be the weights of the portfolio at time t with some method. Then the net returns are given by

$$y_{P,t+1}^{net} = y_{P,t+1} - c(1 + y_{P,t+1}) \sum_{j=1}^{p} |\hat{w}_{t+1,j} - \hat{w}_{t,j} \frac{1 + y_{t+1,j}}{1 + y_{P,t+1}}|,$$

where c is the transaction cost.

Then the out-of-sample mean return is defined as

$$\mu^{net} = \frac{1}{T - T_I} \sum_{t=T_I}^{T-1} y_{P,t+1}^{net}.$$

Note that μ^{net} is the average out-of-sample portfolio return over rolling windows, reported as "Returns" in our tables.

The out-of-sample variance is defined as

$$\hat{\sigma}_{net}^2 = \frac{1}{T - T_I - 1} \sum_{t=T_I}^{T-1} (y_{P,t+1}^{net} - \mu^{net})^2,$$

and reported as "Variance" in our tables.

Lastly, the out-of-sample Sharpe ratio is defined as

$$SR^{net} = \frac{\mu^{net}}{\hat{\sigma}_{net}}.$$

5.1 Data

Our monthly fundamentals dataset covers all S&P 500 constituents from January 2005 to April 2024, sourced from CRSP and Compustat, and includes three firm characteristics: size, book-to-market ratio, and 12-month momentum. This dataset is a subset of Green et al. (2017), and the sample ends in April 2024 to align with the coverage of our news dataset. Following Green et al. (2017), we winsorize all characteristics at the 1st and 99th percentiles, standardize each to zero mean and unit standard deviation, and replace missing values with zero, implying that firms with missing data are assigned the cross-sectional average characteristic for that month.[4]

For each month t's return, we calculate their features at the end of month $t-1$. We assume that annual accounting data are available at this time if the firm's fiscal year ended at least six months before $t-1$.

5.2 Method

5.2.1 Stage 1: Screening Agents

The first stage to our model consists of a screening agent. As described above, this can be a conventional screening agent, such as human analysts, but we also consider LLMs as screening agents. The two such agents are LLM-S, a fundamentals-based LLM that generates buy, sell, and hold signals according to monthly firm fundamentals data, and FinBERT, a news-based agent that generates signals based on sentiment analysis on news each month.

[4]Similar imputation methods for missing data are the standard in empirical asset pricing: used in Gu et al. (2020), Kozak et al. (2020), and Kelly et al. (2025). Chen and McCoy (2024) specifically recommend using mean imputation for ML studies. See footnote 5 in Green et al. (2017) for additional justification.

We first describe the conventional screening approaches. The approach using human analysts is based on analyst recommendations on the IBES dataset. Since analyst recommendations may become stale quickly, we use an exponentially-decreasing weight on each analyst's recommendation, based on how far their date of recommendation is to month-end (we use a half-life of 7 days). We then take the weighted sum of all analysts' recommendation for each S&P500 stock, and calculate the change in the weighted sum month-to-month. If the change in recommendations is greater than 0.5 (respectively, less than -0.5), we interpret that as a sell (respectively, buy) signal. All others are hold signals.

The approach using logistic regression runs a simple logistic regression cross-sectionally for all firms in a specific date. We then evaluate the out-of-sample probabilities for the next month, assign buy signals to firms with probabilities in the top decile, and assign sell signals to firms with probabilities in the bottom decile. Everything else is assigned a hold signal.

Now we describe the LLM screening approaches. FinBERT is a specialized, pre-trained natural language processing model designed for financial sentiment analysis (see Araci (2019)). We use S&P 500 news articles provided online at Hugging Face.[5] For each stock-month combination, we use FinBERT to analyze all news articles in that month. We define a firm's sentiment score to be the positive FinBERT probability minus the negative FinBERT probability and using the same exponentially-decreasing weighted sum as above to account for potentially stale news. If the sentiment score is greater than 0.1 (respectively less than -0.1), we assign a buy signal (respectively assign a sell signal). All other stocks are assigned hold signals.

Lastly, we describe LLM-S. We use Gemini 2.0 Flash in CrewAI and give it the necessary tools to access firm fundamentals data in a given date. It then considers the distribution of firm size, book-to-market, and momentum values for that date, and outputs a deterministic scoring rule for buys, holds, and sell signals. To leverage the zero-shot capabilities of LLMs, we only provide firm characteristics to the LLM at the end of the year, and ask it to provide buy, hold, and sell signals for the next year. Importantly, the agent does not use statistical learning on historical data. Instead, it relies on pretrained domain knowledge and in-context reasoning to analyze the current cross-section of firms and produce screening decisions. Lastly, to ensure a sanity check, we also ask the agent to provide its economic intuition for why it chooses the scoring rule that it does. We have included snapshots of our prompts and an example output in the appendix. To account for any potential look-ahead biases, we explicitly instruct the agent to use causal masking in the prompt.

In the following sections, we also consider ensembles of agents (for example, LLM-S + FinBERT), geared at reducing hallucinations/inaccurate recommendations for each agent. The decision rule for the ensemble is given as follows: given an ensemble of agents A_1 and A_2, who provide sets of buys and sells S_1 and S_2, respectively, the ensemble $A_1 + A_2$ recommends the set $S_1 \cap S_2$ to invest in. However, if the cardinality of the intersection is less than or equal to one stock: $|S_1 \cap S_2| \leq 1$, then it recommends the set $S_1 \cup S_2$. This is saying that ideally, we would like to take

[5] https://huggingface.co/datasets/KrossKinetic/SP500-Financial-News-Articles-Time-Series

the firms in the consensus of both agents. However, if the two agents cannot reach a consensus, we take all of their recommendations. If we have an ensemble of three agents A_1, A_2, and A_3, with recommendations S_1, S_2, and S_3, respectively, then the ensemble $A_1 + A_2 + A_3$ will recommend stocks that belong in at least two distinct sets. This is akin to taking the majority vote between three agents.

The above particular decision rule defaults to the union of sets in case the intersection is trivial. In our analysis in Section 5.3, our best model, FinBERT+LLM-S, has trivial intersections 50% of all dates - hence we default to the union to avoid a mostly empty portfolio. Other ways of defaulting, such as choosing the highest SR single agent (FinBERT) when the intersection is empty, leads to empirically worse performance. We stress that the majority of the performance is due to the intersection rather than the union. For example, in the FinBERT+LLM-S with MV objective and deep learning, 1.037 out of the 1.187 Sharpe ratio is due to the intersection, with the union only having a Sharpe ratio of 0.545. This matches with our intuition, as the intersection is meant towards reducing hallucinations and inaccurate recommendations between the two agents.

We explain our choice of decision rule further. To prevent a hallucination of overselecting stocks and not screening well, we use the intersection rule for two agent system. If their intersection is empty, then it is clear that there is no overselection due to hallucination, and it may be the case that hallucination can go in the direction of underselecting. In that empty intersection scenario, to prevent possible underselecting-hallucination, we use the union of both agents selection. As a robustness check, we also changed our rule in the following way. First we use the intersection rule, and if intersection is empty, we default to the FinBERT agent's selection (Chen et al. (2023) show the strong effect of parsing news sentiments. Our own single agent tables in Tables 2 and 5 also show FinBERT can deliver better performance than LLM-S). This robustness check can be obtained from authors on demand.

5.2.2 Stage 2: Quantitative weighting method

Given a set of buy, sell, and hold signals from Stage 1, we run a variety of statistical techniques and portfolio weight formations to determine the combination that performs best out-of-sample. The statistical techniques used for high-dimensional portfolio formation are nodewise regression, residual nodewise regression, deep learning, and nonlinear shrinkage. The portfolio weight formations are the global minimum variance portfolio, the mean-variance efficient portfolio, and the maximum Sharpe ratio portfolio. All of these methods are described in depth in Section 4.

We only apply these techniques to firms who have either a buy or a sell signal in Stage 1. By considering non-hold firms as a whole, we let the possibility that the quantitative weighting method can correct the LLM screening agent, this can prevent the cascade of errors by LLM-S and FinBERT agents (i.e. assigning a positive weight on a stock that the LLM screening agent had recommended sell to). Lastly, to calculate optimal weights, we utilize a rolling window of 180 months (15 years) of historical returns data, stepping forward one month at a time. Given our

dataset begins in January 2005, the initial 180-month formation period allows our out-of-sample testing to run monthly from January 2020 through April 2024.

5.3 Results

This section presents our empirical results, progressing from baseline comparisons to the evaluation of our full multi-agent framework. We first ask whether pairing an LLM screening agent with a quantitative method improves upon a purely quantitative baseline, and whether any such improvement is specifically attributable to LLM-based screening rather than screening per se. We then examine whether news-based sentiment screening via FinBERT contributes incremental performance gains, and whether integrating traditional analyst judgment with LLM-based screening generates higher Sharpe ratios than either approach alone. Finally, we evaluate whether our full multi-agent AI system delivers superior risk-adjusted performance relative to all single-agent and hybrid alternatives, and whether the best-performing configurations outperform a passive market benchmark. To contextualize these performance metrics, we consistently benchmark our strategies against the S&P 500 index, which generated a Sharpe ratio of 0.6324 between January 2020 and April 2024.

Our main results are in Tables 1-8, which detail annualized Sharpe ratios, returns, and variance metrics. The top Sharpe ratio for each method-portfolio combination is indicated in bold. All our results are after transaction costs of 10 basis points.

5.3.1 Quantitative Method versus LLM-S Plus Quantitative Method

Table 1 presents the baseline case of a purely quantitative strategy without any LLM agent or screening method. In this setup, the NLS/MSR portfolio performs best, achieving a Sharpe ratio of 0.8968, as well as the highest annualized return in the baseline model at 19.99%. Furthermore, operating on the entire universe of S&P 500 stocks allows the baseline to maximally diversify away idiosyncratic risk, achieving the absolute lowest variances across all tables and models (a 0.0180 variance for the baseline POET/MSR portfolio). Notably, this is the only configuration under the purely quantitative baseline that successfully outperforms the broader market's Sharpe ratio of 0.6324; all other quantitative strategies fail to pass this benchmark.

The agent in Table 2 consists of a two-stage model, where the LLM-S agent first generates buy/sell signals, and the quantitative method subsequently determines the portfolio weights. Under this framework, the NLS/GMV portfolio performs best, achieving a Sharpe ratio of 0.6738. When using the LLM-S agent for screening, only the deep learning portfolios and two NLS strategies successfully exceed the market threshold. However, a simple comparison of the maximum Sharpe ratios does not show the full picture. We discuss other portfolio-technique combinations in the next paragraph. As a side note, note that since the quantitative weighting method does not screen, all the buy/sell signals remain constant across all method-portfolio combinations within any given table. The performance variations within each table stem solely from the different precision matrix

estimation techniques and the optimal portfolio construction technique applied to these buy/sell signals.

A critical question is whether incorporating LLM-S screening provides an advantage over the purely quantitative baseline. We see that in all method/portfolio combinations, using LLM-S screening improves all Sharpe ratios (SR henceforth) except for the NLS-POET/MSR portfolios. To illustrate the magnitude of this impact, consider the deep learning/GMV portfolio: without screening, the baseline quantitative method yields an SR of -0.1179 in Table 1, but there is a dramatic improvement to a SR of 0.6607 in Table 2. This illustration demonstrates the effect of LLM-S screening can successfully transform a negative-SR strategy into a profitable one.

5.3.2 Is it screening or LLM-S screening that makes the difference?

To determine whether our observed performance gains are driven by screening in general or uniquely by the LLM-S agent, we establish a benchmark using logistic regression-based screening, detailed in Table 3. Specifically, at each January where we retrain, we use the prior 15 year training window to fit a cross-sectional logistic regression to select the top and bottom decile of firms to long and short, respectively. We then apply the quantitative strategies on the output of the logistic regression to assign portfolio weights. Once again, the NLS/MSR portfolio emerges as the top performer, achieving a Sharpe ratio of 0.7695 in Table 3. Against the market benchmark, only two configurations (the NLS/MSR and deep learning/MSR portfolios) manage to beat the S&P 500 under this screening method. Comparing these results to the LLM-S screening in Table 2, we observe that LLM-S based screening yields higher SRs across all GMV and Markowitz (MV) portfolios. In the MSR portfolios, the evidence is not clear: LLM-S improves the NW portfolio's SR to 0.5425 SR (compared to 0.5182 for logistic screening), though logistic screening seems to be more effective for other precision matrix techniques under the MSR objective.

As a second benchmark, we evaluate screening based on human analyst buy/sell recommendations that are used in conjunction with the quantitative strategies outlined above, with results in Table 4. Relying on human analyst recommendations yields the poorest overall performance, with every single method-portfolio combination failing to beat the market benchmark. Comparing this directly with Table 2 provides a clear assessment of humans versus LLM-S screening capabilities. Across every single method/portfolio combination, the Sharpe ratios generated by LLM-S screening (Table 2) are larger than the ones generated by human analysts (Table 4). This points towards the superiority of the LLM-S model for screening. We attribute this to the over- and under-reaction tendencies well documented in the behavioral finance literature.

We can also address whether any screening method together with a quantitative method improves upon the baseline quantitative strategy (Table 1). Comparing the logistic screening approach (Table 3) with the baseline (Table 1), we see that it does: 13 out of the 15 method-portfolio combinations exhibit improved Sharpe ratios when logistic screening is applied and then these screened stocks are weighted with a quantitative method. But as described above, as a whole, the improve-

ment by using logistic screening is not as large as the gains obtained using LLM-S in GMV-MV portfolios. Conversely, when comparing human analyst screening with the quantitative method (Table 4) to the baseline quantitative strategy, the results are mixed. It is not clear whether human judgment can improve upon the baseline quantitative strategy.

Finally, we benchmark our model against the highly influential screening methodology proposed by Novy-Marx (2013), which utilizes gross profits to assets combined with book-to-market valuation as a screening tool. Novy-Marx (2013) shows that using the profitability of companies, accompanied by the book-to-market ratio, outperforms other possible screening choices such as using only book-to-market. Following the method described in Novy-Marx (2013), out of 500 stocks, we rank the top 150 stocks as "buys" according to the two metric-screening tool described above and designate stocks at the bottom 150 as "sells." Then this screened subset is fed into the quantitative method. The resulting out-of-sample Sharpe ratios are presented in Table C.11. Comparing these results to the LLM-S approach in Table 2, LLM-S maintains an advantage. The highest Sharpe ratio achieved using the Novy-Marx screen is 0.5069 (NLS/MSR), lower than the 0.6738 peak achieved by LLM-S (NLS/GMV). Furthermore, LLM-S provides a higher Sharpe ratio across every method-objective combination. For instance, in the deep learning/MSR portfolio, Novy-Marx screening achieves a 0.4500 SR, while LLM-S screening achieves a SR of 0.6581. This difference is substantial. In terms of returns, as another example, NW/GMV achieves an annual return of 10.6% with LLM-S screening, compared to 6.27% under the Novy-Marx methodology. In Appendix C, we extend this analysis to a 10 year window and analyze the portfolio metrics of the hybrid Novy-Marx/FinBert agent.

In summary, combining LLM-S based screening with quantitative strategies to form the portfolio dominates both the baseline quantitative method and all evaluated conventional screening methods coupled with quantitative methods.

5.3.3 Is FinBERT based screening helpful?

Next, we investigate whether pairing sentiment analysis with the quantitative strategy outperforms the purely quantitative baseline. Comparing Table 1 with Table 5 reveals that screening stocks with FinBERT increases Sharpe ratios across all method-portfolio combinations, with the sole exception of the NLS/MSR portfolio. Furthermore, sentiment-based screening via FinBERT proves highly effective against the broader market: 11 out of 15 method-portfolio combinations beat the S&P 500 Sharpe ratio.

We then assess which AI screening method is more effective: FinBERT or LLM-S? To answer that question, we compare Table 2 with Table 5. The winner in Table 5 (deep learning/GMV) yields a Sharpe ratio of 0.7805 via FinBERT screening, and the winner in Table 2 (NLS/GMV) achieves a Sharpe ratio of 0.6738. While FinBERT achieves a higher maximum Sharpe ratio overall, a direct comparison of each method-portfolio combination yields mixed results between the two models.

Furthermore, we evaluate whether FinBERT-based screening is better than logistical regression-

based screening. To that effect we compare Tables 3 and 5. FinBERT dominates logistical screening across every method-portfolio combination except NLS/MSR. For instance, the deep learning/GMV portfolio achieves a 0.5677 Sharpe ratio under logistic regression screening, but it increases to 0.7805 SR using FinBERT screening - a 37.5% increase. Lastly, comparing human analyst screening (Table 4) versus FinBERT (Table 5) demonstrates that FinBERT uniformly outperforms human judgement. Notably, the advantage FinBERT screening has over human analyst screening is larger than it has over logistic regression screening.

Ultimately, a consistent pattern emerges: FinBERT-based screening outperforms both logistical regression and human-based screening easily, while also substantially improving upon the baseline quantitative strategy as well.

5.3.4 Can Hybrid Screening Help?

In this part, we analyze a hybrid approach that integrates the LLM-S agent with human analyst recommendations. The results of this ensemble are presented in Table 6, and further details can be found in Section 5.2. First, we evaluate this ensemble against the baseline quantitative method by comparing Tables 1 and 6. The results are mixed and there is no clear evidence that incorporating human judgment into the quantitative pipeline yields a performance advantage.

It could be possible that suboptimal decisions by human analysts could have degraded the ensemble's performance. To check that, we compare the hybrid results (Table 6) with the pure LLM-S screening results (Table 2). The highest Sharpe ratio in Table 6 is 0.4968 (NLS/GMV), which falls short of both the 0.6324 market benchmark and the 0.6738 peak Sharpe ratio achieved by pure LLM-S in Table 2. The difference is again moderately large, indicating that human judgment is suppressing the LLM-S Sharpe ratio during this evaluated time period. Furthermore, the results are also uniformly true for every individual method-portfolio comparison. For instance, the deep learning/MSR portfolio achieves a Sharpe ratio of 0.6581 under LLM-S screening, but this decreases to 0.1363 when human analysts are added, a dramatic drop.

5.3.5 Multi-Agent AI versus Single-Agent Screening

We will refer to the ensemble consisting of FinBERT and LLM-S as the Agentic AI system, as it relies on multiple specialized LLM-based agents coordinating with each other and executing different tasks. We want to understand whether there is value to Agentic AI-based screening over the alternative portfolio formations discussed so far. In particular, we evaluate whether the Agentic AI architecture outperforms the baseline quantitative strategy. Then, can Agentic AI do better than standalone single-agent models: LLM-S only or FinBERT only?

In this framework, the agent team operates collaboratively: FinBERT screens based on short-term sentiment analysis, while LLM-S screens based on a monthly fundamentals-driven strategy. Their recommendations are combined together into a consensus signal, which is subsequently processed by the quantitative method to decide the portfolio weights. The results are presented in

Table 7. The results are remarkable when benchmarked against the broader market: all but one of the method-portfolio combinations exceed the S&P 500 Sharpe ratio. Furthermore, the deep learning strategies consistently produce Sharpe ratios above 1. The top-performing configuration is the deep learning/Markowitz portfolio with a 1.1867 Sharpe ratio, representing an 88% improvement over the S&P 500. That is a much larger SR for the same time period compared across all other tables.

We then benchmark the Agentic AI-based approach against the baseline quantitative strategy baseline shown in Table 1. The performance difference is striking. For example, the deep learning/MV portfolio yields a Sharpe ratio of -0.0325. Introducing Agentic AI screening elevates the Sharpe ratio to 1.1867, a remarkable turnaround. Moreover, with the single exception of the NLS/MSR portfolio, the Agentic AI system dominates the baseline across every method-portfolio combination. The differences between the two tables are large; to give another example, the NW/GMV portfolio's Sharpe ratio increases from 0.4506 under the pure quantitative strategy to 0.9212 under the Agentic AI framework, representing a 104% increase.

Beyond Sharpe ratios, the Agentic AI framework generates impressive absolute returns. The peak annualized return is 36.34% in our out-of-sample period, produced by the deep learning/MV portfolio in Table 7. Except for the NLS/MSR portfolio, every method-objective combination in the baseline model generates lower returns compared to its direct counterpart in the Agentic AI ensemble. For instance, the NW/GMV portfolio produces a 7.69% annual return under the quantitative-only strategy, whereas the same method-portfolio combination produces a 22.01% return annually with the Agentic AI framework—nearly a three-fold increase.

These returns/variance patterns illustrate a fundamental *predictability-diversification* tradeoff. The purely quantitative strategy operates on the entire universe of the S&P 500 (no screening), which allows it to diversify away idiosyncratic risk and achieve the absolute lowest variances out of all tables and models, such as the 0.0180 variance by the POET/MSR portfolio. In contrast, the Agentic AI framework acts as a high conviction screener, investing in an average of 22 stocks based on strong fundamentals and positive news sentiment. By holding such concentrated portfolios, the Agentic AI models naturally sacrifice some diversification benefits, resulting in higher variances. Notably, the lowest variance among all Agentic AI portfolios is 0.0387, achieved by the deep learning/MSR portfolio, which already exceeds that of 10 out of 15 quantitative baseline portfolios. In addition, its highest-returning portfolio (deep learning/MV) carries a large variance of 0.0938. However, the high predictability and significant excess returns generated by the Agentic AI portfolios far outweigh the diversification penalty. For instance, while the highly diversified baseline POET/MSR portfolio minimizes variance (0.0180), it only yields an annualized return of 5.71% and a corresponding Sharpe ratio of 0.4262. On the other hand, the Agentic AI deep learning/MV portfolio results in a higher variance of 0.0938, but compensates with a 36.34% annualized return, pushing its Sharpe ratio to an impressive 1.1867. All this demonstrates why the Agentic AI system delivers superior performance: the predictability that Agentic AI demonstrates outweighs

the diversification benefits it sacrifices by taking a screened portfolio of stocks.

Next, we address whether a single agent by itself (FinBERT or LLM-S) can deliver better results than Agentic AI. To evaluate this, we compare the Agentic AI results (Table 7) against the FinBERT-only (Table 5) and LLM-S only (Table 2) frameworks. The evidence heavily favors Agentic AI. The FinBERT-only maximum Sharpe ratio is 0.7805 (deep learning/GMV), compared to the Agentic AI winner of 1.1867 (deep learning/Markowitz). Across all individual method-portfolio combination, the Agentic AI system significantly outperforms the FinBERT-only model, with the sole exception of the NLS/MSR portfolio. For instance, the deep learning/GMV portfolio has a Sharpe ratio of 0.7805 in Table 5, but this increases to 1.0148 when using Agentic AI.

Next, we analyze the performance of Agentic AI compared to LLM-S in Tables 2 and 7. The highest Sharpe ratio achieved by LLM-S is 0.6738, compared to 1.1867 with Agentic AI. Across every method-portfolio combination, Agentic AI consistently dominates LLM-S. For example, LLM-S delivers a Sharpe ratio of 0.5911 in the NW/GMV portfolio, and this increases to 0.9212 when using Agentic AI. The increase is large and can make a large difference in practice.

The main reason why Agentic AI outperforms single agents is that it selects firms with desirable fundamentals *and* news, with the intersection as an opportunity for the two separate agents to mitigate each other's errors. To demonstrate this, we consider a deliberately suboptimal version of Agentic AI where the system screens for stocks with desirable news, but not with strong fundamentals. On our best Agentic AI method-portfolio combination (deep learning/MV), the Sharpe ratio drops to 0.5523 from 1.1867. Similarly, when Agentic AI screens for stocks with strong fundamentals but weak news, the Sharpe ratio drops to 0.6662. In addition, this is evidence that dropping certain stocks from the portfolio can lead to a significant increase in the out-of-sample Sharpe ratio. This justifies our decision to screen for stocks first, and our Agentic AI model chooses those stocks with promising long-term fundamentals and short-term news, making this a natural choice for a screening agent.

In summary, there is a substantial gain by using Agentic AI. It can easily surpass single agent systems, and even though we omitted a detailed comparison versus human analysts or logistical regression screening, the differences are large and can be seen by comparing Tables 3 and 4 with Table 7. The results can further be interpreted as the deviation from a benchmark with large number of assets. Our multi-agent AI ensemble selects a substantially lower number of stocks out of a large universe[6], and hence takes on a large tracking error, compared with S&P 500, to achieve larger returns and Sharpe ratios. Theoretical results on tracking error in portfolios have been recently established in high dimensions in Caner and Fan (2026).

[6]To be specific, our Agentic AI model selected 22 stocks on average out of 500 in our out-sample period from January 2020 to April 2024.

5.3.6 Multi-Agent AI with Human Analyst Input

Finally, we investigate whether integrating human judgment into the Agentic AI framework yields any benefits. To evaluate this, Table 8 presents the results of an ensemble consisting of FinBERT, LLM-S, and human analysts working together. We benchmark this against the pure Agentic AI system detailed in Table 7. It is clear that adding human judgment significantly degrades performance, reducing the maximum SR from 1.1867 (deep learning/Markowitz) to 0.5941 (NW/GMV). Furthermore, in this three-agent ensemble, all resulting Sharpe ratios fall below the market benchmark, underscoring the detrimental effect of human judgment. This degradation is also consistently true when we compare each method-portfolio individually. To give a particularly stark example, the deep learning/MSR portfolio achieves a Sharpe ratio of 1.0107 when using Agentic AI, but this decreases to -0.0108 when human analysts are added.

5.4 Larger Time Span

In this section, we analyze a larger time span between January 2015-April 2024 (spanning 9 years and 4 months). This extended period includes the former time span between January 2020-April 2024, but introduce more diverse market environments, including the strong bull markets of 2017 and 2019. Consequently, the S&P 500 index exhibits a higher baseline Sharpe ratio of 0.7298 over this duration. We now revisit the core questions outlined in Subsections 5.3.1-5.3.6 under this expanded time-frame.

The corresponding results are in Tables C.1 to C.8. First, we evaluate whether pairing the LLM-S agent with the quantitative strategy (Table C.2) improves upon the purely quantitative baseline (Table C.1). Comparing the highest Sharpe ratios achieved within each table suggests that LLM-S screening maximum SR actually declines from 0.7187 (POET/MSR) to 0.6367(NLS/GMV). However, a more detailed look shows that deep learning and residual nodewise techniques benefit from LLM-S screening, demonstrating improved Sharpe ratios relative to their baseline counterparts.

When we benchmark against logistic regression-based screening (Table C.3), the deep learning/MSR portfolio surprisingly emerges as the winner with a Sharpe ratio of 0.8531. This not only improves upon its respective baseline-Table C.1 but also exceeds the winners when using LLM-S screening. However, we see that there is some promise in the NLS technique, as LLM-based screening is better than logistical screening for two of three portfolio construction techniques. We also compare logistic screening against sentiment-based FinBERT screening (Table C.5). FinBERT screening provides uniformly better Sharpe ratios compared to logistic screening. In both cases, the deep learning/MSR portfolio is the winner with 0.8817 and 0.8531 in Tables C.5 and C.3 respectively.

Screening driven by human analyst recommendations (Table C.4) continues to underperform. The best-performing human analyst model is still lower than the winner of the baseline model in Table C.1. Across all method-portfolio combinations, human analyst screening is worse compared to the baseline quantitative-only method, with deep learning portfolios and residual nodewise-MSR

portfolios being the exception.

Conversely, sentiment screening via FinBERT paired with the quantitative method (Table C.5) yields impressive results, delivering much higher Sharpe ratios uniformly over the baseline in Table C.1. To give an example, the deep learning/MSR portfolio SR rebounds from a -0.2314 baseline SR to a profitable 0.8817 SR. That is a remarkable turnaround.

Consistent with our findings from the 5-year sample, incorporating human analyst recommendations alongside AI agents (Tables C.6 and C.8) perform poorly. These tables show poor SR results compared with Table C.1. The winners in these tables yield much lower Sharpe ratios compared to the winner in the unscreened baseline model.

Finally, we address our central question: does the Agentic AI framework provide an advantage over an extended time horizon of just under 10 years? The results from Table C.7 confirm that it does. The Agentic AI ensemble achieves the highest overall SR of 0.9429 via the deep learning/MSR portfolio, easily surpassing the peak SRs in all other tables. To highlight the magnitude of this impact, the same deep learning/MSR configuration yields a SR of -0.2314 in the baseline model (Table C.1), but surges to a SR of 0.9429 under the Agentic AI architecture (Table C.7). Furthermore, this portfolio is also much higher than the S&P 500 SR of 0.7298. Also we compare our Table C.7 with LLM-S with quantitative weighting in Table C.2. SR of GMV and MSR portfolios in LLM-S quantitative weighting are dominated by SR in LLM-S+ FinBERT and quantitative weights (Agentic AI). To give an example, GMV-NLS portfolio has SR of 0.6367 as the best one in Table C.2, and the NLS-GMV portfolio SR is 0.8290. The difference is very large. Then we compare our Table C.7 with Table C.5 to see that whether Agentic AI makes a difference compared to only the FinBert agent plus the quantitative strategy. The results are mixed: the Agentic AI SRs are higher for GMV portfolios but lower with MV portfolio. But the top SR in Table C.7 (0.9429) still dominates the top SR in Table C.5 (0.8817).

BASELINE-ONLY WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.4506	0.4742	0.4769	0.0769	0.0792	0.0798	0.0291	0.0279	0.0280
Residual NW	0.0464	0.0651	-0.1030	0.0066	0.0089	-0.0213	0.0201	0.0187	0.0428
Deep learning	-0.1179	-0.0325	-0.4972	-0.0377	-0.0115	-0.5752	0.1024	0.1253	1.3381
POET	0.2961	0.3538	0.4262	0.0411	0.0477	0.0571	0.0193	0.0182	0.0180
NLS	0.3978	0.3964	**0.8968**	0.0560	0.0550	0.1999	0.0198	0.0192	0.0497

Table 1: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to all firms in the S&P500. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.5911	0.6079	0.5425	0.1060	0.1028	0.0943	0.0321	0.0286	0.0302
Residual NW	0.6026	0.5979	0.2501	0.0835	0.0819	0.0472	0.0192	0.0187	0.0356
Deep learning	0.6607	0.6736	0.6581	0.1025	0.1033	0.0998	0.0241	0.0235	0.0230
POET	0.4716	0.5234	0.4060	0.0755	0.0813	0.0689	0.0257	0.0241	0.0288
NLS	**0.6738**	0.6549	0.5091	0.0968	0.0939	0.0941	0.0206	0.0205	0.0342

Table 2: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that the LLM has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

6 Conclusion

This paper introduces a multi-agent Agentic AI platform for portfolio management. Our architecture coordinates three specialized agents: an LLM-Strategy agent (LLM-S) that screens stocks annually on the basis of fundamental firm characteristics, a FinBERT sentiment agent that screens monthly on the basis of financial news, and a quantitative weighting method that applies high-dimensional precision matrix estimation techniques to determine optimal portfolio weights over the screened asset universe. We have shown that the full system, evaluated on the S&P 500 over 2020–2024, achieves a peak annualized Sharpe ratio of 1.1867, an 88% improvement over the market, while generating peak annualized returns of 36.34%.

We have shown that performance gains depends on the system's design. Adding human analyst recommendations to the AI ensemble consistently degrades performance. This underscores that Agentic AI is not simply a powerful substitute for human judgment, rather it is a qualitatively different conceptual framework, one that avoids the behavioral and emotional biases that systematically compromise human financial decision-making.

LOGISTIC REGRESSION WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.4684	0.4722	0.5182	0.0799	0.0788	0.0855	0.0291	0.0278	0.0272
Residual NW	0.4170	0.4646	0.4156	0.0583	0.0646	0.1048	0.0195	0.0194	0.0637
Deep learning	0.5677	0.5592	0.6924	0.0916	0.0889	0.1135	0.0261	0.0253	0.0269
POET	0.4265	0.4503	0.4456	0.0636	0.0660	0.0703	0.0222	0.0215	0.0249
NLS	0.5602	0.4995	**0.7695**	0.0828	0.0736	0.2055	0.0219	0.0217	0.0713

Table 3: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that logistic regression has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

HUMAN ANALYSTS WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.4238	0.3344	**0.4331**	0.0747	0.0571	0.0736	0.0310	0.0292	0.0289
Residual NW	-0.1324	-0.1893	-0.0442	-0.0186	-0.0265	-0.0099	0.0197	0.0195	0.0498
Deep learning	0.3285	0.2257	0.3658	0.0515	0.0353	0.0571	0.0245	0.0244	0.0243
POET	0.2475	0.1097	0.2473	0.0397	0.0176	0.0400	0.0258	0.0257	0.0262
NLS	0.1226	0.0284	0.2294	0.0172	0.0040	0.0570	0.0197	0.0200	0.0618

Table 4: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that analysts have screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

FINBERT WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.6797	0.5789	0.6406	0.1351	0.1268	0.1218	0.0395	0.0480	0.0361
Residual NW	0.6463	0.5432	0.6087	0.1125	0.1137	0.1020	0.0303	0.0438	0.0281
Deep learning	**0.7805**	0.6378	0.7582	0.1429	0.1291	0.1299	0.0335	0.0409	0.0294
POET	0.7029	0.6255	0.7133	0.1453	0.1388	0.1414	0.0427	0.0492	0.0393
NLS	0.7386	0.6349	0.6502	0.1249	0.1268	0.1120	0.0286	0.0399	0.0297

Table 5: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that FinBERT have screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S + HUMAN ANALYSTS WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.4732	0.3686	0.2450	0.0963	0.0794	0.0495	0.0414	0.0463	0.0409
Residual NW	0.4299	0.3672	-0.1618	0.0799	0.0711	-0.0382	0.0345	0.0375	0.0558
Deep learning	0.4634	0.3752	0.1363	0.0893	0.0764	0.0277	0.0371	0.0415	0.0414
POET	0.4539	0.4562	0.2763	0.0952	0.0995	0.0584	0.0440	0.0476	0.0446
NLS	**0.4968**	0.4412	-0.0635	0.0955	0.0855	-0.0155	0.0370	0.0376	0.0595

Table 6: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that LLM+analysts have screened, from IBES. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S + FINBERT WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.9212	1.1232	0.8564	0.2201	0.3481	0.1836	0.0571	0.0960	0.0459
Residual NW	0.9366	1.1137	0.6885	0.2010	0.3397	0.1446	0.0461	0.0930	0.0441
Deep learning	1.0148	**1.1867**	1.0107	0.2288	0.3634	0.1988	0.0508	0.0938	0.0387
POET	0.9339	1.1673	0.9035	0.2117	0.3402	0.1878	0.0514	0.0850	0.0432
NLS	0.8745	1.0368	0.5884	0.1943	0.3265	0.1279	0.0494	0.0992	0.0473

Table 7: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that FinBERT+LLM has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S + FINBERT + HUMAN ANALYSTS WITH QUANTITATIVE WEIGHTING: 2020-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	**0.5941**	0.3286	0.1900	0.1197	0.0667	0.0379	0.0406	0.0412	0.0398
Residual NW	0.4259	0.2454	-0.5689	0.0740	0.0426	-0.1329	0.0302	0.0302	0.0546
Deep learning	0.5812	0.3268	-0.0108	0.1090	0.0634	-0.0021	0.0352	0.0376	0.0388
POET	0.5669	0.3391	0.2098	0.1148	0.0698	0.0434	0.0410	0.0424	0.0428
NLS	0.4659	0.2981	-0.4574	0.0866	0.0539	-0.1124	0.0346	0.0327	0.0604

Table 8: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that LLM+FinBERT+analysts have screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

We have shown that the intersection-based decision rule, under which a stock is included in the candidate portfolio only when both LLM-S and FinBERT agree, is responsible for the majority of the performance gain. In the deep learning/MV configuration, for instance, 1.037 of the 1.187 Sharpe ratio is attributable to the intersection, with the union fallback contributing only 0.545. This demonstrates that the hallucination-suppression role of multi-agent consensus is a significant empirical driver of returns. The result also highlights the economic mechanism at work: by requiring agreement between an agent that screens on long-run fundamentals and one that screens on short-run sentiment, the system selects firms that are simultaneously undervalued and positively perceived by the market, combining two complementary sources of alpha.

From a theoretical point of view, we make a contribution to the high-dimensional portfolio literature. The number of assets in our portfolio is a random variable, realized through the screening process before any weights are estimated. We have shown that under sensible screening, the squared Sharpe ratio of the screened portfolio consistently estimates its target, even under mild screening errors. The result holds for any precision matrix estimator satisfying standard consistency conditions, and is therefore broadly applicable across the quantitative portfolio formation literature.

Acknowledgments

The authors received no specific funding for this work. Furthermore, the authors declare that they have no financial relationships or other potential conflicts of interest.

Appendix

A Proofs

Investing is a two choice model, both the number of stocks as well as weights are key to this choice. When we screen the universe p of stocks, our choice via screening will be a random variable $\hat{p}$. In other words, we estimate a bounded random variable $1 \leq \hat{p} \leq p$ via screening. One screening method we advocate here is Agentic AI. Screening amounts to choosing the names of stocks, but as a simplification for our consistency analysis, this can be done in a sensible way as defined in Section 3.1.

The aim of the screening is to select a certain stocks which may be optimal according to a strategy as in momentum based strategy or a metric like Sharpe Ratio. The optimal number of stocks in this strategy or a metric is defined as, non-random, p^*, $1 \leq p^* \leq p$. If there is an optimal number (non-random) of stocks p^*, we want to see that if our screening process is allowing mild mistakes, our estimated Sharpe Ratio with these screened stocks can achieve this target Sharpe Ratio of the optimal p^* with probability approaching one. This is a new idea to see whether a mildly successful screening can achieve Sharpe Ratio consistency. We use the following Assumption A.1 that allows mild mistake in selecting stocks.

.

Assumption A.1. With $\hat{p}, p^*$ growing with n

$$\frac{|\hat{p} - p^*|}{p^*} = o_p(1).$$

Since $\hat{p}, p^*$ are positive integers and grow with n, this assumption amounts to a slight mistake of a constant difference between $\hat{p}, p^*$. Assumption A.1 is used instead of the restrictive consistent integer estimation, which is perfect selection of the number of stocks in a portfolio

$$\lim_{n\to\infty} P\left[\left|\frac{\hat{p}}{p^*} = 1\right|\right] = 1.$$

Even though a mistake can be made in selecting the correct number of stocks (i.e. hence stock indexes), this may not affect downstream Sharpe Ratio estimation consistency. This is a novel result and a new proof. We show this in Lemma A.1-A.2, and the proof of Theorem 1 below.

Next we show that precision matrix, $\Gamma_{p^*} : p^* \times p^*$ can be estimated consistently by $\hat{\Gamma}_p^* : p^* \times p^*$ at a given p^*. Define the maximum row sum of a matrix $A : m \times n$ as $\|A\|_{l_\infty} = \max_{1\leq i\leq m} \sum_{j=1}^{n} |A_{i,j}|$, where $A_{i,j}$ represents (i,j) th element of matrix A.

Assumption A.2.

$$\|\hat{\Gamma}_{p^*} - \Gamma_{p^*}\|_{l_\infty} = o_p(1).$$

For example both Callot et al. (2021), Caner et al. (2023) prove Assumption A.2 under weaker assumptions. Spectral norm consistency such as in Fan et al. (2013), Caner and Daniele (2025) can also be used, and this will change the proofs but not the result of consistency we conjecture.

Next we analyze Global Minimum Variance portfolio and specifically its Sharpe Ratio (SR from now on). The aim of the portfolio is to minimize the variance of the portfolio of assets. We will show that if we do screen stocks, the high dimensional consistency of SR will be obtained. Denote the precision matrix estimator after screening as $\hat{\Gamma}_{\hat{p}} : \hat{p} \times \hat{p}$, and the mean return estimator as $\hat{\mu}_{\hat{p}} : \hat{p} \times 1$. These estimators can come from any technique. Estimated Sharpe Ratio of Global Minimum Variance Portfolio is

$$\widehat{SR}_{\hat{p}} := \sqrt{\hat{p}} \left(\frac{1'_{\hat{p}} \hat{\Gamma}_{\hat{p}} \hat{\mu}_{\hat{p}}}{\hat{p}} \right) \left(\frac{1'_{\hat{p}} \hat{\Gamma}_{\hat{p}} 1_{\hat{p}}}{\hat{p}} \right)^{-1/2},$$

which is defined in (27) of Caner et al. (2023) with non-random p instead of $\hat{p}$. In Caner et al. (2023), as well as all the literature we know of, the estimated SR is using a fixed-non-random large universe of assets p. Here since there is a screening process before estimating SR, the dimension of the chosen portfolio is $\hat{p} \leq p$.

Denote the target precision matrix as $\Gamma_{p^*} : p^* \times p^*$, and the mean return is $\mu_{p^*} : p^* \times 1$. The target Sharpe Ratio is:

$$SR_{p^*} := \sqrt{p^*} \left(\frac{1'_{p^*} \Gamma_{p^*} \mu_{p^*}}{p^*} \right) \left(\frac{1'_{p^*} \Gamma_{p^*} 1_{p^*}}{p^*} \right)^{-1/2},$$

which is defined in (26) of Caner et al. (2023) but with p the universe of assets, rather than target screened number of assets p^*. In Caner et al. (2023), as well as the other literature such as Callot et al. (2021),Caner and Daniele (2025), target SR is for p, universe of assets. Here we use the target number of screened stocks, p^*.

The main technical complication is even though screening process can allow mistakes, it is not clear a major financial metric such as Sharpe Ratio can be consistently estimated in high dimensions.

In order to show the consistency of Sharpe Ratio estimate, we first need the following Lemma. It will show consistency of variance estimator for Global Minimum Variance portfolio. We have two possibilities, either random variable $\hat{p}$ is larger than or equal to target p^*, or smaller than equal to the target, p^*. For each possibility, we need a norm constraint on blocks of precision matrix estimates.

First, it is possible to overshoot the target, implying $\hat{p} \geq p^*$. Decompose $\hat{\Gamma}_{\hat{p}} : \hat{p} \times \hat{p}$ into four blocks

$$\hat{\Gamma}_{\hat{p}} = \begin{bmatrix} \hat{\Gamma}_{p^*} & \hat{\Gamma}_{p^*,1} \\ \hat{\Gamma}_{p^*,2} & \hat{\Gamma}_{p^*,3} \end{bmatrix}, \tag{A.1}$$

where $\hat{\Gamma}_{p^*} : p^* \times p^*$, $\hat{\Gamma}_{p^*,1} : p^* \times \hat{p} - p^*$, $\hat{\Gamma}_{p*,2} : \hat{p} - p^* \times p^*$, $\hat{\Gamma}_{p^*,3} : \hat{p} - p^* \times \hat{p} - p^*$.

Other possibility is we may undershoot the target $\hat{p} \leq p^*$,and for that we decompose the precision matrix estimator $\hat{\Gamma}_{p^*} : p^* \times p^*$ as

$$\hat{\Gamma}_{p^*} = \begin{bmatrix} \hat{\Gamma}_{\hat{p}} & \hat{\Gamma}_{\hat{p},1} \\ \hat{\Gamma}_{\hat{p},2} & \hat{\Gamma}_{\hat{p},3} \end{bmatrix}. \tag{A.2}$$

Block dimensions are: $\hat{\Gamma}_{\hat{p}} : \hat{p} \times \hat{p}$, $\hat{\Gamma}_{\hat{p},1} : \hat{p} \times (p^* - \hat{p})$, $\hat{\Gamma}_{\hat{p},2} : (p^* - \hat{p}) \times \hat{p}$, and $\hat{\Gamma}_{\hat{p},3} : (p^* - \hat{p}) \times (p^* - \hat{p})$.

Condition A.1 imposes that the blocks that emanate from over or undershooting in screening can have maximum row sums bounded by a constant. This is a reasonable condition imposed since either row or columns are small in number. However, we should note that either with high $\hat{p}$ or p^*, this constraint may become more strict. This condition can be relaxed to allow diverging sums if we know the rate of convergence of $\hat{p}$ to p^* in consistency.

Condition A.1. *When (i).$\hat{p} \geq p^*$ we need $\|\hat{\Gamma}_{p^*,m}\|_{l_\infty} \leq C < \infty$, $m = 1, 2, 3$. When (ii). $\hat{p} \leq p^*$ we need $\|\hat{\Gamma}_{\hat{p},m}\|_{l_\infty} \leq C < \infty$,$m = 1, 2, 3$.*

Note that Condition A.1 can be relaxed in the following way. For $\hat{p} \geq p^*$ we can have $\|\hat{\Gamma}_{p^*,m}\|_{l_\infty} \leq C < \infty$ with probability approaching one, and same for (ii) too. The proofs will not change.
Now we introduce following Lemma A.1 that shows GMV variance consistency.

Lemma A.1. *Under Assumptions A.1-A.2 with Condition A.1 we have*

$$\frac{|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}1_{p^*}|}{p^*} = o_p(1).$$

Proof of Lemma A.1

We start the proof by the analysis of

$$|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}1_{p^*}|. \tag{A.3}$$

.

Step 1. To consider (A.3) we start with $\hat{p} \geq p^*$ case. In that scenario we can decompose the vector of ones and the precision matrix estimator as follows

$$1_{\hat{p}} = \begin{pmatrix} 1_{p^*} \\ 1_{\hat{p}-p^*} \end{pmatrix}, \quad \hat{\Gamma}_{\hat{p}} = \begin{bmatrix} \hat{\Gamma}_{p^*} & \hat{\Gamma}_{p^*,1} \\ \hat{\Gamma}_{p^*,2} & \hat{\Gamma}_{p^*,3} \end{bmatrix} \tag{A.4}$$

where vector of ones is decomposed into $p^*, \hat{p} - p^*$, parts $1_{p^*}, 1_{\hat{p}-p^*}$ respectively. Similarly the precision matrix estimator is decomposed into four blocks, explained in (A.1). Then rewrite (A.3)

by adding and subtracting $1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*}$

$$\begin{aligned} |1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}1_{p^*}| &= |1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} - 1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*} + 1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*} - 1'_{p^*}\Gamma_{p^*}1_{p^*}| \\ &\leq |1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} - 1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*}| + |1'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})1_{p^*}|. \end{aligned} \tag{A.5}$$

Consider the analysis of the first term on the right side of (A.5), use (A.1)(A.4) and triangle inequality

$$\begin{aligned} |1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} \quad - \quad 1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*}| &\leq |1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,2}1_{p^*}| \\ + \quad |1'_{p^*}\hat{\Gamma}_{p^*,1}1_{\hat{p}-p^*}| + |1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,3}1_{\hat{p}-p^*}|. \end{aligned} \tag{A.6}$$

Analyze each term in (A.6). Start with the first right side term

$$\begin{aligned} |1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,2}1_{p^*}| &\leq \|1'_{\hat{p}-p^*}\|_1\|\hat{\Gamma}_{p^*,2}1_{p^*}\|_\infty \\ &\leq \|1_{\hat{p}-p^*}\|_1\|\hat{\Gamma}_{p^*,2}\|_{l_\infty}\|1_{p^*}\|_\infty \\ &= |\hat{p} - p^*|\|\hat{\Gamma}_{p^*,2}\|_{l_\infty}, \end{aligned} \tag{A.7}$$

where we use Holder's inequality for the first inequality, and for the second inequality we use p.345 of Horn and Johnson (2013).

Start with the second right side term, by taking a transpose of that term first in absolute terms

$$\begin{aligned} |1'_{\hat{p}-p^*}\hat{\Gamma}'_{p^*,1}1_{p^*}| &\leq \|1'_{\hat{p}-p^*}\hat{\Gamma}'_{p^*,1}\|_1\|1_{p^*}\|_\infty \\ &\leq \|1_{\hat{p}-p^*}\|_1\|\hat{\Gamma}_{p^*,1}\|_{l_\infty} \\ &\leq |\hat{p} - p^*|\|\hat{\Gamma}_{p^*,2}\|_{l_\infty}, \end{aligned} \tag{A.8}$$

where we use Holder's inequality for the first inequality, and for the second inequality we use p.345 of Horn and Johnson (2013), and $\|A'\|_{l_1} = \|A\|_{l_\infty}$, for a generic matrix A.

The analysis for the third term right side term in (A.6) is very same as in the first right side term analysis in (A.7) , so we have

$$\frac{1}{p^*}|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} - 1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*}| = o_p(1), \tag{A.9}$$

given Assumption A.1, and by Condition A.1. For the second term on the right side of (A.5).

$$\begin{aligned} |1'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})1_{p^*}| &\leq \|1'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})\|_\infty\|1_{p^*}\|_1 \\ &= \|1_{p^*}\|_\infty\|\hat{\Gamma}_{p^*} - \Gamma_{p^*}\|_{l_\infty}p^*, \end{aligned} \tag{A.10}$$

where we use Holder's inequality for the first inequality, and for the second inequality we use p.345

of Horn and Johnson (2013). Then use Assumption A.2 to have

$$\left| \frac{1'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})1_{p^*}}{p^*} \right| = o_p(1). \tag{A.11}$$

Combine (A.9)(A.11) in (A.5) to have

$$\left| \frac{1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}1_{p^*}}{p^*} \right| = o_p(1). \tag{A.12}$$

Step 2. Now we consider the opposite case, $\hat{p} \leq p^*$. The indicator can be decomposed as

$$1_{p^*} = \begin{pmatrix} 1_{\hat{p}} \\ 1_{p^*-\hat{p}} \end{pmatrix}, \tag{A.13}$$

Note that the vector ones decomposed into two subvectors of dimension $\hat{p}, p^* - \hat{p}$ respectively. By adding and subtracting $1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*}$, we begin with a triangle inequality

$$|1'_{p^*}\Gamma_{p^*}1_{p^*} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}}| \leq |1'_{p^*}\Gamma_{p^*}1_{p^*} - 1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*}| + |1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}}|. \tag{A.14}$$

We simplify the second term on the right side in (A.14) using (A.2)(A.13)

$$\begin{aligned} |1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}}| &= |1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}} + 1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},2}1_{\hat{p}} + 1'_{\hat{p}}\hat{\Gamma}_{\hat{p},1}1_{p^*-\hat{p}} + 1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},3}1_{p^*-\hat{p}} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}}| \\ &\leq |1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},2}1_{\hat{p}}| + |1'_{\hat{p}}\hat{\Gamma}_{\hat{p},1}1_{p^*-\hat{p}}| + |1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},3}1_{p^*-\hat{p}}|. \end{aligned} \tag{A.15}$$

where the first and fifth elements cancel in the equality in (A.15).
In (A.15) consider the first right side term in the last inequality

$$\begin{aligned} |1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},2}1_{\hat{p}}| &\leq \|1_{p^*-\hat{p}}\|_1 \|\hat{\Gamma}_{\hat{p},2}1_{\hat{p}}\|_\infty \\ &\leq \|1_{p^*-\hat{p}}\|_1 \|\hat{\Gamma}_{\hat{p},2}\|_{l_\infty} \\ &\leq (p^* - \hat{p})\|\hat{\Gamma}_{\hat{p},2}\|_{l_\infty}, \end{aligned} \tag{A.16}$$

where we use Holder's inequality for the first inequality, and p.345 of Horn and Johnson (2013) for the second inequality. The second term on the right side of (A.15) is transposed

$$\begin{aligned} |1'_{p^*-\hat{p}}\hat{\Gamma}'_{\hat{p},1}1_{\hat{p}}| &\leq \|1'_{p^*-\hat{p}}\hat{\Gamma}'_{\hat{p},1}\|_1 \|1_{\hat{p}}\|_\infty \\ &\leq \|1'_{p^*-\hat{p}}\|_1 \|\hat{\Gamma}'_{\hat{p},1}\|_{l_1} \\ &= (p^* - \hat{p})\|\hat{\Gamma}_{\hat{p},1}\|_{l_\infty}, \end{aligned} \tag{A.17}$$

where the first inequality is by Holder's inequality, and the second one is by p.345 of Horn and Johnson (2013), and the equality is by $\|A'\|_{l_1} = \|A\|_{l_\infty}$. The third right side term in the last inequality in (A.15) is handled in the same way as in (A.16). So via Assumption A.1 with the Condition A.1 about block matrices $\|\hat{\Gamma}_{\hat{p},m}\|_{l_\infty} \leq C < \infty$, $m = 1, 2, 3$ the scaled second term on the right side of (A.14) is

$$\frac{|1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}}|}{p^*} = o_p(1). \tag{A.18}$$

Next, use (A.11) on the first right side term in (A.14)

$$\frac{|1'_{p^*}\Gamma_{p^*}1_{p^*} - 1'_{p^*}\hat{\Gamma}_{p^*}1_{p^*}|}{p^*} = o_p(1). \tag{A.19}$$

Combine the last two equations on the left side term of (A.14)

$$\frac{|1'_{p^*}\Gamma_{p^*}1_{p^*} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}}|}{p^*} = o_p(1). \tag{A.20}$$

Step 3. Combine Steps 1-2 (A.12)(A.20) to have the desired result. **Q.E.D.**

We provide the following assumptions for the next lemma.

Assumption A.3.

(i).

$$\max_{1\leq j\leq p} |\mu_j| \leq C < \infty.$$

(ii).

$$\max_{1\leq j\leq p} \|\hat{\mu}_{\hat{p}} - \mu_p\|_\infty = o_p(1).$$

This assumption is for all the universe of assets p, which is non-random. Note that Assumption A.3(i) is standard in the literature and can be seen in Caner et al. (2023). Assumption A.3(ii) can be proven under weaker conditions and a rate of convergence can be obtained. This is shown in Caner et al. (2023), and also in this assumption we let $\hat{\mu}_p$ any consistent estimator of μ_p. Next assumption is a norm bound on the target precision matrix.

Assumption A.4.

$$\|\Gamma_{p^*}\|_{l_\infty} \leq C < \infty.$$

Assumption A.4 puts a finite bound on maximum row sums of a $p^* \times p^*$ target precision matrix. In case of p^* growing with n assumption is restrictive. But in that scenario, for a specific estimator like residual nodewise regression of Caner et al. (2023) the proof may depend on the joint product of estimation of error mean estimator as in Assumption A.3(ii) (but with p^* dimension, $p^* \leq p$)

multiplied by the target precision matrix in Assumption A.4. Then we can have weaker assumptions that shows divergence of the sum in Assumption A.4, but multiplied by the rate in Assumption A.3(ii) can converge to zero in probability. In other words, it is possible to have

$$\max_{1\leq j\leq p^*} \|\hat{\mu}_{p^*} - \mu_{p^*}\|_\infty = o_p(d_n), \quad d_n \to 0,$$

$$\|\Gamma_{p^*}\|_{l_\infty} = O(r_n), \quad r_n \to \infty$$

but $d_n r_n \to 0$. For details, see Sharpe Ratio estimation for non-random, no screening GMV portfolio of Caner et al. (2023) proof. Since we have a general estimation setup, we impose more strict assumptions. Specific estimators can relax these assumptions.

Lemma A.2. *Under Assumptions A.1-A.4 with Condition A.1*

$$\frac{\left|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}\mu_{p^*}\right|}{p^*} = o_p(1).$$

Proof of Lemma A.2.

We start with the numerator by adding and subtracting $1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*}$ and triangle inequality

$$|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}\mu_{p^*}| \leq |1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} - 1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*}| + |1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*} - 1'_{p^*}\Gamma_{p^*}\mu_{p^*}|. \tag{A.21}$$

In the next two steps, first we consider $\hat{p} \geq p^*$, and then $\hat{p} \leq p^*$.

Step 1. We start with case $\hat{p} \geq p^*$.

We can decompose the estimator of the mean vector

$$\hat{\mu}_{\hat{p}} = \begin{bmatrix} \hat{\mu}_{p^*} \\ \hat{\mu}_{\hat{p}-p^*} \end{bmatrix},$$

where $\hat{\mu}_{p^*} : p^* \times 1$, and $\hat{\mu}_{\hat{p}-p^*} : \hat{p} - p^* \times 1$. Now substitute this decomposition immediately above, and for $\hat{\Gamma}_{\hat{p}}$ in Step 1 in Lemma A.1 proof, using (A.1)(A.4)

$$1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} = 1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*} + 1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,2}\hat{\mu}_{p^*} + 1'_{p^*}\hat{\Gamma}_{p^*,1}\hat{\mu}_{\hat{p}-p^*} + 1_{\hat{p}-p^*}\hat{\Gamma}_{p^*,3}\hat{\mu}_{\hat{p}-p^*}. \tag{A.22}$$

Consider the first term on the right side of (A.21), with (A.22) and triangle inequality

$$\begin{aligned} |1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} - 1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*}| \quad &\leq \quad |1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,2}\hat{\mu}_{p^*}| \\ &+ \quad |1'_{p^*}\hat{\Gamma}_{p^*,1}\hat{\mu}_{\hat{p}-p^*}| + |1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,3}\hat{\mu}_{\hat{p}-p^*}|. \end{aligned} \tag{A.23}$$

Before the next proof,

$$\|\hat{\mu}_{p^*}\|_\infty \leq \|\hat{\mu}_{\hat{p}}\|_\infty \leq \|\hat{\mu}_{\hat{p}} - \mu_{\hat{p}}\|_\infty + \|\mu_{\hat{p}}\|_\infty = o_p(1) + O(1), \tag{A.24}$$

where we use $p^* \leq p$ and Assumption A.3.

Analyze the first term on the right side of (A.23)

$$\begin{aligned} |1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,2}\hat{\mu}_{p^*}| &\leq \|1_{\hat{p}-p^*}\|_1\|\hat{\Gamma}_{p^*,2}\hat{\mu}_{p^*}\|_\infty \\ &\leq (\hat{p}-p^*)\|\hat{\Gamma}_{p^*,2}\|_{l_\infty}\|\hat{\mu}_{p^*}\|_\infty \\ &= (\hat{p}-p^*)O(1)o_p(1), \end{aligned} \tag{A.25}$$

where we use Holder's inequality for the first inequality, and p.345 of Horn and Johnson (2013) for the second inequality, and the rates come from Condition A.1, and (A.24). Then via Assumption A.1

$$\frac{|1'_{\hat{p}-p^*}\hat{\Gamma}_{p^*,2}\hat{\mu}_{p^*}|}{p^*} = o_p(1). \tag{A.26}$$

Next, we consider second term on the right side of (A.23)

$$\begin{aligned} |1'_{p^*}\hat{\Gamma}_{p^*,1}\hat{\mu}_{\hat{p}-p^*}| &= |\hat{\mu}'_{\hat{p}-p^*}\hat{\Gamma}'_{p^*,1}1'_{p^*}| \\ &\leq \|\hat{\mu}'_{\hat{p}-p^*}\hat{\Gamma}'_{p^*,1}\|_1\|1_{p^*}\|_\infty \\ &\leq \|\hat{\mu}_{\hat{p}-p^*}\|_1\|\hat{\Gamma}'_{p^*,1}\|_{l_1} = \|\hat{\mu}_{\hat{p}-p^*}\|_1\|\hat{\Gamma}_{p^*,1}\|_{l_\infty} \\ &\leq (\hat{p}-p^*)[\max_{1\leq j\leq \hat{p}-p^*}|\hat{\mu}_j|]\|\hat{\Gamma}_{p^*,1}\|_{l_\infty} \\ &\leq (\hat{p}-p^*)[\max_{1\leq j\leq p}|\hat{\mu}_j|]\|\hat{\Gamma}_{p^*,1}\|_{l_\infty} \\ &= (\hat{p}-p^*)O_p(1)O(1), \end{aligned} \tag{A.27}$$

where we use Holder's inequality for the first inequality, and p.345 of Horn and Johnson (2013) for the second inequality, and for the second equality we use $\|A'\|_{l_1} = \|A\|_{l_\infty}$, and for the third inequality we use $l_1 - l_\infty$ inequality ($\|v\|_1 \leq dim(v)\|v\|_\infty$, with dim(v) showing the dimension of vector v), and the fourth inequality by $\hat{p} - p^* \leq p$, and the rates are by Condition A.1 and (A.24). By (A.27) and Assumption A.1

$$\frac{|1'_{p^*}\hat{\Gamma}_{p^*,1}\hat{\mu}_{\hat{p}-p^*}|}{p^*} = o_p(1). \tag{A.28}$$

Third term on the right side of (A.23) is considered in the same way as in (A.28), so combining (A.26)(A.28) with third term analysis we have in left side term of (A.23)

$$\frac{|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} - 1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*}|}{p^*} = o_p(1). \tag{A.29}$$

Now consider the second term on the right side of (A.21). By adding and subtracting $1'_{p^*}\Gamma_{p^*}\hat{\mu}_{p^*}$ via triangle inequality

$$|1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*} - 1'_{p^*}\Gamma_{p^*}\mu_{p^*}| \leq |1'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})\hat{\mu}_{p^*}| + |1'_{p^*}\Gamma_{p^*}(\hat{\mu}_{p^*} - \mu_{p^*})|. \tag{A.30}$$

First, we consider the first right side term in (A.30)

$$\begin{aligned} |1'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})\hat{\mu}_{p^*}| &= |\hat{\mu}'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})'1_{p^*}| \\ &\leq \|\hat{\mu}'_{p^*}(\hat{\Gamma}_{p^*} - \Gamma_{p^*})'\|_1 \|1_{p^*}\|_\infty \\ &\leq \|\hat{\mu}_{p^*}\|_1 \|(\hat{\Gamma}_{p^*} - \Gamma_{p^*})'\|_{l_1} \\ &\leq p^*[\max_{1\leq j\leq p^*} |\hat{\mu}_j|] \|\hat{\Gamma}_{p^*} - \Gamma_{p^*}\|_{l_\infty} \\ &= p^* O_p(1) o_p(1), \end{aligned} \tag{A.31}$$

where we use Holder's inequality for the first inequality, and p.345 of Horn and Johnson (2013) for the second inequality, and for the third inequality we use $\|A'\|_{l_1} = \|A\|_{l_\infty}$, and $l_1 - l_\infty$ vector norm inequality, and for the rates we use the same analysis in (A.24) with Assumption A.2.

Next, consider the second term on the right side of (A.30)

$$\begin{aligned} |1'_{p^*}\Gamma_{p^*}(\hat{\mu}_{p^*} - \mu_{p^*})| &\leq \|1'_{p^*}\Gamma_{p^*}\|_1 \|\hat{\mu}_{p^*} - \mu_{p^*}\|_\infty \\ &\leq \|1_{p^*}\|_1 \|\Gamma_{p^*}\|_{l_1} \|\hat{\mu}_{p^*} - \mu_{p^*}\|_\infty \\ &= p^* \|\Gamma_{p^*}\|_{l_\infty} \|\hat{\mu}_{p^*} - \mu_{p^*}\|_\infty \\ &= p^* O(1) o_p(1), \end{aligned} \tag{A.32}$$

where we use Holder's inequality for the first inequality, and p.345 of Horn and Johnson (2013) for the second inequality, and the equality is by symmetricity of Γ, and the rate is by Assumptions A.3-A.4 and $p^* \leq p$. Combine (A.31)(A.32) into (A.30)

$$\frac{|1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*} - 1'_{p^*}\Gamma_{p^*}\mu_{p^*}|}{p^*} = o_p(1). \tag{A.33}$$

Combine (A.29)(A.33) into (A.21) to have

$$\frac{|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}\mu_{p^*}|}{p^*} = o_p(1). \tag{A.34}$$

Step 2.

Consider (A.21) but the case of $p^* \geq \hat{p}$. In that scenario decompose the mean estimator

$$\hat{\mu}_{p^*} := \begin{bmatrix} \hat{\mu}_{\hat{p}} \\ \hat{\mu}_{p^*-\hat{p}} \end{bmatrix},$$

where $\hat{\mu}_{\hat{p}} : \hat{p} \times 1$ and $\hat{\mu}_{p^*-\hat{p}} : p^* - \hat{p} \times 1$ dimension. We see that by (A.13)

$$\begin{aligned} 1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*} &= (1'_{\hat{p}}, 1'_{p^*-\hat{p}}) \begin{bmatrix} \hat{\Gamma}_{\hat{p}}, \hat{\Gamma}_{\hat{p},1} \\ \hat{\Gamma}_{\hat{p},2}, \hat{\Gamma}_{\hat{p},3} \end{bmatrix} \begin{bmatrix} \hat{\mu}_{\hat{p}} \\ \hat{\mu}_{p^*-\hat{p}} \end{bmatrix} \\ &= 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} + 1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},2}\hat{\mu}_{\hat{p}} \\ &+ 1'_{\hat{p}}\hat{\Gamma}_{\hat{p},1}\hat{\mu}_{p^*-\hat{p}} + 1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},3}\hat{\mu}_{p^*-\hat{p}}, \end{aligned} \tag{A.35}$$

where we use $\hat{\Gamma}_{p^*}$ decomposition before Condition A.1. Then the first term on the right side of (A.21) simplifies via (A.35)

$$\begin{aligned} |1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}}| &= |1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},2}\hat{\mu}_{\hat{p}} + 1'_{\hat{p}}\hat{\Gamma}_{\hat{p},1}\hat{\mu}_{p^*-\hat{p}} + 1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},3}\hat{\mu}_{p^*-\hat{p}}| \\ &\leq |1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},2}\hat{\mu}_{\hat{p}}| + |1'_{\hat{p}}\hat{\Gamma}_{\hat{p},1}\hat{\mu}_{p^*-\hat{p}}| \\ &+ |1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},3}\hat{\mu}_{p^*-\hat{p}}|. \end{aligned} \tag{A.36}$$

See that

$$\|\hat{\mu}_{\hat{p}}\|_\infty = \max_{1\leq j\leq \hat{p}} |\hat{\mu}_j| \leq \max_{1\leq j\leq p} |\hat{\mu}_j| \leq \max_{1\leq j\leq p} |\mu_j| + \max_{1\leq j\leq p} |\hat{\mu}_j - \mu_j| = O_p(1), \tag{A.37}$$

by $\hat{p} \leq p^* \leq p$ and Assumption A.3 via triangle inequality in the last inequality. Also we have by the same analysis in (A.37)

$$\|\hat{\mu}_{p^*-\hat{p}}\|_\infty = \max_{1\leq j\leq p^*-\hat{p}} |\hat{\mu}_j| = O_p(1). \tag{A.38}$$

Next, consider the first term on the the right side of (A.36).

$$\begin{aligned} |1'_{p^*-\hat{p}}\hat{\Gamma}_{\hat{p},2}\hat{\mu}_{\hat{p}}| &= |\hat{\mu}'_{\hat{p}}\hat{\Gamma}'_{\hat{p},2}1_{p^*-\hat{p}}| \leq \|\hat{\mu}_{\hat{p}}\|_\infty \|\hat{\Gamma}'_{\hat{p},2}1_{p^*-\hat{p}}\|_1 \\ &\leq \|\hat{\mu}_{\hat{p}}\|_\infty \|\hat{\Gamma}'_{\hat{p},2}\|_{l_1}\|1_{p^*-\hat{p}}\|_1 \\ &= \|\hat{\mu}_{\hat{p}}\|_\infty \|\hat{\Gamma}_{\hat{p},2}\|_{l_\infty}(p^* - \hat{p}) \\ &= O_p(1)O(1)(p^* - \hat{p}), \end{aligned} \tag{A.39}$$

where the first inequality is by Holder's inequality, and the second inequality is by p.345 of Horn and Johnson (2013), and the second equality is by $\|A'\|_{l_1} = \|A\|_{l_\infty}$ for a generic matrix A, and the rates are by (A.37), and Condition A.1.

Now we consider the second term on the right side of (A.36)

$$\begin{aligned} |1'_{\hat{p}}\hat{\Gamma}_{\hat{p},1}\hat{\mu}_{p^*-\hat{p}}| &\leq \|1'_{\hat{p}}\hat{\Gamma}_{\hat{p},1}\hat{\mu}_{p^*-\hat{p}}\|_\infty \|\hat{\mu}_{p^*-\hat{p}}\|_1 \leq (p^* - \hat{p})\|1_{p^*}\|_\infty \|\hat{\Gamma}_{\hat{p},1}\|_{l_\infty}\|\hat{\mu}_{p^*-\hat{p}}\|_\infty \\ &= (p^* - \hat{p})O(1)O_p(1), \end{aligned} \tag{A.40}$$

where we use Holder's inequality for the first inequality, and p.345 of Horn and Johnson (2013) and $l_1 - l_\infty$ vector norm inequality for the second inequality, and the rates are from Condition A.1

and (A.38). The third term on the right side of (A.36) is handled in the same way as in (A.40). Combine (A.39)(A.40) into the left side term in (A.36) (which is also the first right side term in (A.21)

$$\frac{|1'_{p^*}\hat{\Gamma}_{p^*}\hat{\mu}_{p^*} - 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}}|}{p^*} = o_p(1), \tag{A.41}$$

by Assumption A.1. Now second right term in (A.21) is handled in the same way as in Step 1 proof- (A.33). So via (A.33)(A.41) we have

$$\frac{|1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}} - 1'_{p^*}\Gamma_{p^*}\mu_{p^*}|}{p^*} = o_p(1). \tag{A.42}$$

Step 3. Steps 1-2 both show the same result so Lemma A.2 is proved. **Q.E.D**

We need the following assumption, with $Eigmin(A)$ representing the minimum eigenvalue for a generic square matrix A, and $c > 0$ is a positive constant.

Assumption A.5.

(i).

$$\frac{|1'_{p^*}\Gamma_{p^*}\mu_{p^*}|}{p^*} \geq c > 0$$

(ii).

$$Eigmin(\Gamma_{p^*}) \geq c > 0$$

Both assumptions are used in the literature, for (i) see Caner et al. (2023), Caner and Fan (2026), and for (ii), see Fan et al. (2011), Callot et al. (2021). If we consider first , the numerator in Assumption (i), this is a sum of returns, μ_{p^*}, scaled by variance, Γ_{p^*} since p^* may grow with n, this sum may grow with p^*, so we scale with p^* as well. We want the scaled returns to be larger than a positive constant. Assumption (ii) can be proved in factor models as in Fan et al. (2011), Caner et al. (2023), Caner and Fan (2026).

Theorem A.1

Under Assumptions A.1-A.5

$$\left| \frac{\widehat{SR}^2_{\hat{p}}}{SR^2_{p^*}} - 1 \right| = o_p(1).$$

Remarks. 1. This is a new result in high dimensional Sharpe Ratio estimation for a GMV portfolio. The number of stocks to be selected is treated as a random variable $\hat{p}$, and selected via a screening method. It is possible to have mild mistakes, $\hat{p} \neq p^*$, in this portfolio selection. After that selection, we show that this Sharpe ratio via the screened assets in the portfolio can be consistent estimator for a target Sharpe Ratio with a mildly different selection of stocks p^*. The literature before can only show high dimensional consistency of Sharpe Ratio with fixed number of stocks, no screening process is used, or screened stocks are treated as a nonrandom quantity. For this literature with

unconstrained an constrained high dimensional portfolios see Caner et al. (2023), Caner and Fan (2026).

2. We still allow $p > n$ consistency for screened Sharpe Ratio estimation as long as the precision matrix-mean return are estimated consistently in the case of $p > n$ as in Fan et al. (2011), Callot et al. (2021) Caner et al. (2023).

Proof of Theorem A.1. Note that by Assumption A.5

$$1'_{p^*}\Gamma_{p^*}1^*_p \geq \|1_{p^*}\|_2^2 Eigmin(\Gamma_{p^*}) \geq cp^*. \tag{A.43}$$

Then we try to simplify the notation in the proof by defining $\hat{y} := 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}\hat{\mu}_{\hat{p}}$, $y = 1'_{p^*}\Gamma_{p^*}\mu_{p^*}$, $\hat{x} = 1'_{\hat{p}}\hat{\Gamma}_{\hat{p}}1_{\hat{p}}$, $x := 1'_{p^*}\Gamma_{p^*}1_{p^*}$.

Then squared Sharpe Ratio estimation can be written as

$$\begin{aligned} \frac{\widehat{SR_{\hat{p}}}^2}{SR^2_{p^*}} - 1 &= \frac{\hat{p}(\hat{y}^2/\hat{x})}{p^*(y^2/x)} - 1 \\ &= \frac{\hat{p}}{p^*}\left(\frac{\hat{y}^2}{y^2}\right)\left(\frac{x}{\hat{x}}\right) - 1 \end{aligned} \tag{A.44}$$

We can rewrite the right side of (A.44)

$$\begin{aligned} \frac{\widehat{SR}^2_{\hat{p}}}{SR^2_{p^*}} - 1 &= \frac{\hat{p}}{p^*}\left(\frac{\hat{y}^2 - y^2 + y^2}{y^2}\right)\left(\frac{x}{\hat{x} - x + x}\right) - 1 \\ &= \frac{\hat{p}}{p^*}\left(\frac{\hat{y}^2 - y^2}{y^2} + 1\right)\left(\frac{x}{\hat{x} - x + x}\right) - 1 \end{aligned} \tag{A.45}$$

We consider first

$$\begin{aligned} \left|\frac{\hat{y}^2 - y^2}{y^2}\right| &\leq \frac{|\hat{y} - y||\hat{y} + y|}{y^2} \\ &\leq \left[\left|\frac{\hat{y} - y}{y}\right|\right]\left[\left|\frac{\hat{y} - y}{y}\right| + 2\right] \\ &= \left[\left|\frac{\hat{y} - y}{y}\right|\right]^2 + 2\left[\left|\frac{\hat{y} - y}{y}\right|\right] \\ &= \left[\frac{p^* o_p(1)}{cp^*}\right]^2 + 2\left[\frac{p^* o_p(1)}{cp^*}\right] = o_p(1), \end{aligned} \tag{A.46}$$

where for the rate we use Lemma A.2 and Assumption A.5(i) with $y, \hat{y}$ definitions. Next, consider the following term above in (A.45), with $x, \hat{x}$ definitions,

$$\frac{x}{\hat{x}-x+x} \leq \left|\frac{x}{x-|\hat{x}-x|}\right| = \left|\frac{1}{1-\frac{|\hat{x}-x|}{x}}\right| \leq \frac{1}{1-o_p(1)} \leq 1+o_p(1), \tag{A.47}$$

where the result is derived from (A.43), Lemma A.1. Use (A.45)(A.46) on the right side of (A.45) to have

$$\begin{aligned}\left|\frac{\widehat{SR}^2_{\hat{p}}}{SR^2_{p^*}}-1\right| &\leq \frac{\hat{p}}{p^*}(1+o_p(1))(1+o_p(1))-1 \\ &\leq \left|\frac{\hat{p}-p^*+p^*}{p^*}\right|(1+o_p(1))(1+o_p(1))-1 \\ &= (1+o_p(1))(1+o_p(1))(1+o_p(1))-1=o_p(1). \end{aligned} \tag{A.48}$$

Q.E.D.

B Additional Details on Precision matrix estimation techniques

In this section, we describe the precision matrix estimation techniques we use in detail.

B.1 Nodewise regression

Nodewise regression, first introduced by Meinshausen and Bühlmann (2006) and applied to risk-estimation in portfolio settings by Callot et al. (2021), estimates the precision matrix directly via lasso type penalized linear regressions. Meinshausen and Bühlmann (2006) starts with precision matrix formulation. In an high dimensional world, to achieve consistency, by imposing sparsity through zero elements in rows of the precision matrix, they use a penalized estimation to achieve that sparsity in the estimators. Following the framework in Callot et al. (2021), to obtain each row of the precision matrix estimate, we consider p separate regressions.

Let y_t be the $p \times 1$ vector of all asset excess returns at time t. Nodewise regression runs p different Lasso regressions to obtain the precision matrix. In this method, there is exact sparsity assumption on the rows of the precision matrix. This sparsity is defined by the maximum number of nonzero elements across rows of the precision matrix as $\bar{s}$, and this number should be smaller than n. Since there is sparsity on the precision matrix, the estimator should match that. One of the best ways to impose sparsity on estimators is through lasso, as suggested by Meinshausen and Bühlmann (2006).

Let $y_{t,j}$ represent excess asset return of asset j. Let $y_{t,-j}$ represent all asset returns (excess) apart from asset j, and this is a $p-1 \times 1$ vector. For each asset $j = 1, \cdots, p$, we model its returns $y_{t,j}$ related to other assets' returns $y_{t,-j}$ through the following (assuming all returns are time-demeaned):

$$y_{t,j} = y'_{t,-j}\gamma_j + \eta_{t,j},$$

where γ_j are the regression coefficients and $\eta_{t,j}$ is the error term. Note that sparsity assumption on the elements of each row of precision matrix is imposed on slope γ_j basically. This is clear from Step 3 of our Algorithm below. Hence, we will run a lasso regression and penalizing the elements of γ_j.

The estimation algorithm can be described as follows:

1. **Lasso regression**. To account for high dimensions, the coefficients γ_j above are estimated via Lasso regression; that is:

$$\hat{\gamma}_j = \text{argmin}_\gamma \left[\frac{||y_j - Y_{-j}\gamma||_2^2}{n} + 2\lambda_j ||\gamma||_1 \right].$$

   Here, $\hat{\gamma}_j$ is a vector of length $p-1$ that are estimates of γ_j, and λ_j is a positive tuning parameter that determines the size of the l_1 penalty of the estimates.

2. **Parameter selection**. The tuning parameter λ_j is chosen to minimize the GIC criterion across a set of possible λ_j choices-a grid search- (see Fan and Tang (2013)), defined as

$$GIC(\lambda_j) := \log(\hat{\sigma}^2_{\lambda_j}) + |\hat{S}_{\lambda_j}| \frac{\log p}{n} \log(\log n),$$

   where $\hat{\sigma}^2_{\lambda_j} = \frac{||y_j - Y_{-j}\hat{\gamma}_j||_2^2}{n}$ is the mean squared error of the Lasso regression and $|\hat{S}_{\lambda_j}|$ is the cardinality of the set of nonzero parameters in $\hat{\gamma}_j$ using λ_j. Fan and Tang (2013) show that the GIC selects the true model with probability approaching 1 both when $p < n$ and $p \geq n$. The optimal λ_j is denoted as λ_j^*, and the optimal slope is $\hat{\gamma}_j(\lambda_j^*)$, which means the lasso regression which uses λ_j^* among the grid search in Step 1 provides the optimal slope estimate. For ease of notation denote $\hat{\gamma}_j := \hat{\gamma}_j(\lambda_j^*)$.

3. **Precision matrix construction**. Repeat steps 1 and 2 for all $j = 1, \ldots, p$. As derived in Callot et al. (2021), then the precision matrix can be constructed directly by using matrix algebra and defining the diagonal elements $\hat{\Gamma}_{j,j} = \hat{\tau}_j^{-2}$ and the vector of off-diagonal elements for $j-th$ row $\hat{\Gamma}_{j,-j} = -\hat{\tau}_j^{-2}\hat{\gamma}'_j$, where $\hat{\tau}_j^2 = \frac{||y_j - Y_{-j}\hat{\gamma}_j||_2^2}{n} + \lambda_j ||\hat{\gamma}_j||_1$. We form each row j of $\hat{\Gamma}$ by using the main diagonal term $\hat{\Gamma}_{j,j}$ and the off-diagonal term in row j as $\hat{\Gamma}_{j,-j}$. Then stacking each row one upon other we form $\hat{\Gamma}$. $\hat{\Gamma}$ is then the nodewise estimator of the precision matrix.

Under certain assumptions (see Callot et al. (2021) and Chang et al. (2018)), the nodewise regression estimate $\hat{\Gamma}$ consistently estimates the precision matrix even when $p > n$.

B.2 Residual nodewise regression

Proposed by Caner et al. (2023), residual nodewise regression is an extension on nodewise regression by integrating factor models. Unlike nodewise regression, which produces a sparse precision matrix of returns, residual nodewise regression only assumes a sparse precision matrix of the idiosyncratic errors, allowing for a dense precision matrix of returns.

The asset returns are modeled by $y_{t,j} = b_j' f_t + u_{t,j}$, where $f_t : K \times 1$ are observable factors. In our empirical analysis, we use the standard Fama-French three factor model. But in theory we allow growing number of factors. Here, $b_j : K \times 1$ are the factor loadings for K factors, and $u_{t,j}$ are unobserved errors. The precision matrix of the returns, $\Gamma = \Sigma^{-1}$, can be estimated by the following steps:

1. **Factor Removal**. Firstly, we estimate the residuals $\hat{u}_{t,j}$ via OLS. That is,

$$\hat{u}_j = y_j - X'\hat{b}_j,$$

   where $X = (f_1, \ldots, f_n) : K \times n$ is the matrix of the factors and the factor loadings are $\hat{b}_j = (XX')^{-1}Xy_j$, the OLS estimator. $y_j : n \times 1$ vector of returns for asset j, $y_j := (y_{1,j}, \cdots, y_{t,j}, \cdots, y_{n,j})'$. In matrix form, let us write the matrix of factor loadings as $\hat{B} = (YX')(XX')^{-1}$, where $Y : p \times n$ matrix, where each row represents the asset returns, and the columns represent time periods.

2. **Nodewise on Residuals**. Next, we apply nodewise regression, described in Section 4.1.1, on the residuals $\hat{u}$ to estimate $\hat{\Omega}$ which will be the estimate for Σ_u^{-1}, the precision matrix of the errors.

3. **Reconstruction**. We reconstruct the final precision matrix via the Sherman-Morrison-Woodbury formula:

$$\hat{\Gamma} = \hat{\Omega} - \hat{\Omega}\hat{B}[\hat{\Sigma}_f^{-1} + \hat{B}'\hat{\Omega}_{sym}\hat{B}]^{-1}\hat{B}'\hat{\Omega},$$

   where $\hat{\Omega}_{sym} = \frac{\hat{\Omega}+\hat{\Omega}'}{2}$ is the symmetric version of $\hat{\Omega}$, $\hat{B}$ is the matrix of estimated factor loadings, and $\hat{\Sigma}_f = \frac{1}{n}XX' - \frac{1}{n^2}X1_n1_n'X'$ is the sample covariance of the factors.

Likewise, under certain assumptions, there is a consistency guarantee of the precision matrix of asset returns when $p > n$, established in Caner et al. (2023). However, the main difference is in the number of factors. Compared to nodewise regression, convergence is slower for factor models with many factors, and is the reason we choose a factor model with a small number of factors.

B.3 Principal Orthogonal Complement Thresholding (POET)

Proposed by Fan et al. (2013), the principal orthogonal complement thresholding method (POET) is a powerful method to estimate the covariance matrix for linear factor when variables share

common factors but are unobservable. The asset returns are modeled by a linear factor model $y_{t,j} = b_j' f_t + u_{t,j}$, however, the factors at time t, $f_t : K \times 1$ are unknown and need to be estimated, $b_j : K \times 1$ factor loadings for asset j. They assume that covariance matrix of errors is sparse, see (2.2) of Fan et al. (2013). POET first uses principal components analysis (PCA) to estimate the unobservable factors, then uses a thresholding method to estimate the covariance matrix of errors. The detailed steps can be describe as follows:

1. **Factor Estimation**. First, estimate the unobserved $F : n \times K$ matrix of factors by principal components. To that effect, denote $Y : p \times n$ matrix of asset returns. The largest K eigenvalues of $n \times n$ matrix $Y'Y$ is found and the K eigenvectors corresponding to them designated as the columns of $\hat{F}/\sqrt{n}$. The estimate of factor loadings is $\hat{B} = n^{-1} Y \hat{F} : p \times K$.

2. **Number of Factors**. To estimate the number of unknown factors, use Bai and Ng (2002) formula, let $\hat{F}_k : n \times k$ matrix of estimated factors. Let $\|A\|_F$ be the Frobenius norm of matrix A.
$$\hat{K} = \text{argmin}_{1 \le k \le K_1} \left[\log(\frac{1}{pn} \|Y - \frac{1}{n} Y \hat{F}_k \hat{F}_k'\|_F^2) + \frac{k(p+n)}{pn} \log(\min\{p, n\}) \right],$$
where K_1 is a constant upper bound.

3. **Residual Estimation**. Set the residuals $\hat{u}_{jt} = y_{t,j} - \hat{b}_j' \hat{f}_t$ where $\hat{f}_t : \hat{K} \times 1$ is the t th row of $\hat{F}$, (in column vector form) and $\hat{b}_j'$ is the j th row of $\hat{B}$. Then estimate the components of the sample covariance matrix of errors $\hat{\Sigma}_u = \left(\hat{\sigma}_{ij}\right)$ by
$$\hat{\sigma}_{ij} = \frac{1}{n} \sum_{t=1}^{n} \hat{u}_{t,i} \hat{u}_{t,j}.$$

4. **Thresholding**. For each component $\hat{\sigma}_{ij}$, compute a threshold
$$\hat{\tau}_{ij} = \frac{1}{2} \left(\frac{1}{\sqrt{p}} + \sqrt{\frac{\log(p)}{n}} \right) \sqrt{\frac{1}{n} \sum_{t=1}^{n} (\hat{u}_{it} \hat{u}_{jt} - \hat{\sigma}_{ij})^2}.$$
Denote the thresholded covariance matrix as $\hat{\Sigma}_{u,Th}$. If $\hat{\sigma}_{ij} < \hat{\tau}_{ij}$, set the (i,j)-th component of $\hat{\Sigma}_u$ to be 0, otherwise, keep $\hat{\sigma}_{ij}$. Its invertibility is proved in Fan et al. (2013).

5. **Precision Matrix Construction**. We reconstruct the final precision matrix by:
$$\hat{\Gamma} = \hat{\Sigma}_{u,Th}^{-1} - \hat{\Sigma}_{u,Th}^{-1} \hat{B} (I_K + \hat{B}' \hat{\Sigma}_{u,Th}^{-1} \hat{B})^{-1} \hat{B}' \hat{\Sigma}_{u,Th}^{-1}.$$

POET is high dimensional consistent under pervasive factor assumption, with $p > n$ case too.

B.4 Deep learning

When the asset returns have non-linear relationships with factors, estimating the correlation between different assets becomes challenging. Inspired by Farrell et al. (2021), Caner and Daniele (2025) introduced a deep learning-based method to estimate the precision matrix for non-linear factor models.

The asset returns are modeled by $y_{t,j} = g_j(f_t) + u_{t,j}$, where f_t is a K-dimension observable column vector and $g_j(\cdot)$ is an unknown function. Then the covariance matrix of asset returns Σ_y can be decomposed to $\Sigma_g + \Sigma_u$, where Σ_g representing the covariance of unknown function $g_j(.)$, and Σ_u represents the covariance matrix of errors. In order to estimate the precision matrix, Caner and Daniele (2025) use a multi-layer neural network to capture the non-linear relationships.

1. **Non-linear Factor estimation**. To estimate $g_j(\cdot)$, fit f_t and $y_{t,j}$ into a deep neural network. Then for the j-th asset and feature f_t, we are able to get the prediction $\hat{g}_j$. The estimation of Σ_g is
$$\hat{\Sigma}_g = \frac{1}{n}\sum_{t=1}^{n}(\hat{g}_j(f_t) - \bar{g}_j(f_t))(\hat{g}_j(f_t) - \bar{g}_j(f_t))',$$
where $\bar{g}_j(f_t) = \frac{1}{n}\sum_{t=1}^{n}\hat{g}_j(f_t)$.

2. **Residual Estimation**. Set the deep learning residuals $\hat{u}_{t,j} = y_{t,j} - \hat{g}_j(f_t)$. Then estimate the components of the sample covariance matrix of errors $\hat{\Sigma}_u = \left(\hat{\sigma}_{ij}\right)$ by
$$\hat{\sigma}_{ij} = \frac{1}{n}\sum_{t=1}^{n}\hat{u}_{t,i}\hat{u}_{t,j}$$

3. **Thresholding**. For each component $\hat{\sigma}_{ij}$, compute a threshold
$$\hat{\tau}_{ij} = Cr_n\sum_{t=1}^{n}|\hat{u}_{t,i}\hat{u}_{t,j} - \hat{\sigma}_{ij}|,$$
where C is a positive constant, $r_n^{1/2} = n^{-\beta/2(\beta+K)}(\log n)^4$ and β is a smoothness parameter of $g_j(\cdot)$. Note that function estimation error by deep learning estimator is r_n. Denote the thresholded covariance matrix as $\hat{\Sigma}_{u,Th}$. If $\hat{\sigma}_{ij} < \hat{\tau}_{ij}$, set the (i,j)-th component of $\hat{\Sigma}_{u,Th}$ to be 0, otherwise, keep $\hat{\sigma}_{ij}$.

4. **Precision Matrix Construction**. We reconstruct the final precision matrix by:
$$\hat{\Gamma} = \hat{\Sigma}_{u,Th}^{-1} - \hat{\Sigma}_{u,Th}^{-1}\hat{\Sigma}_g(I_k + \hat{\Sigma}_{u,Th}^{-1}\hat{\Sigma}_g)^{-1}\hat{\Sigma}_{u,Th}^{-1}.$$

The thresholding method here is different from POET due to a deep learning-nonlinear model. Under several assumptions on $g_j(\cdot)$, the eigenvalues of Σ_g and Σ_u, deep learning estimator of the

precision matrix is consistent but only when the number of assets, p is smaller than the time period, n.

B.5 Nonlinear Shrinkage (NLS)

Introduced by Ledoit and Wolf (2017) and advanced by Ledoit and Wolf (2020), nonlinear shrinkage address the problem of sample covariance matrix instability in high dimensions through modifying its eigenvalues. Compared to earlier approaches, this analytical method derives a closed-form solution using the Hilbert transform. Let $Y : n \times p$ matrix of excess asset returns, where the rows represent time series and the columns represent various assets.

1. Start with the sample covariance $S := Y'Y/n$. It admits a spectral decomposition $S := U\Lambda U'$, where $U : p \times p$ orthogonal matrix, that has eigenvectors columns of U. Obtain U. Let Λ be the diagonal matrix of eigenvalues of S.

2. We then provide Hilbert transform estimators based on Epanechnikov kernel that will be used in shrinkage function estimator. There will be three cases of interest. First, when $p < n$, all $\lambda_j > 0$, and when $p > n$ there are two possibilities, either $\lambda_j > 0$, or $\lambda_j = 0$, since $p > n$, sample covariance matrix is singular.

   (i). For $\lambda_j > 0$, $p < n$ case, with definition of adaptive bandwidth adjusted for sample eigenvalues, $h_{n,k} := n^{-1/3}\lambda_k$, $k = 1, \cdots, p$

$$\begin{aligned}\hat{H}_f(\lambda_j) &:= \frac{1}{p}\sum_{k=1}^{p}\{\frac{-3(\lambda_j - \lambda_k)}{10\pi h_{n,k}^2} \\ &+ \frac{3}{(4\sqrt{5})\pi h_{n,k}}\left[1 - \frac{1}{5}\left(\frac{(\lambda_j - \lambda_k)}{h_{n,k}}\right)^2\right] log\left|\frac{\sqrt{5}h_{n,k} - \lambda_j + \lambda_k}{\sqrt{5}h_{n,k} + \lambda_j - \lambda_k}\right|\}\end{aligned}$$

   (ii). In the case of $\lambda_j > 0, p > n$ case

$$\begin{aligned}\hat{H}_{\bar{f}}(\lambda_j) &:= \frac{1}{n}\sum_{k=p-n+1}^{p}\{\frac{-3(\lambda_j - \lambda_k)}{10\pi h_{n,k}^2} \\ &+ \frac{3}{(4\sqrt{5})\pi h_{n,k}}\left[1 - \frac{1}{5}\left(\frac{(\lambda_j - \lambda_k)}{h_{n,k}}\right)^2\right] log\left|\frac{\sqrt{5}h_{n,k} - \lambda_j + \lambda_k}{\sqrt{5}h_{n,k} + \lambda_j - \lambda_k}\right|\}\end{aligned}$$

(iii). In case of $\lambda_j = 0, p > n$ case

$$\begin{aligned}\hat{H}_{\bar{f}}(0) &= \left[\frac{3}{10h_n^2} + \frac{3}{4\sqrt{5}h_n}\left(1 - \frac{1}{5h_n^2}\right) log\left(\frac{1+\sqrt{5}h_n}{1-\sqrt{5}h_n}\right)\right] \\ &\times \frac{1}{\pi n}\sum_{j=p-n+1}^{p}\frac{1}{\lambda_j}.\end{aligned}$$

Section 4.7 of Ledoit and Wolf (2020) uses $h_n := n^{-1/3}$ as the bandwidth.

3. In the case of $\lambda_j > 0$, for the optimal shrinkage function estimation in the next step we need an estimate of spectral density function. First for the case with $p < n$, the spectral density, we use Epanechnikov kernel

$$\hat{f}(\lambda_j) := \frac{1}{p}\sum_{k=1}^{p}\frac{3}{4\sqrt{5}h_{n,k}}\left[1 - \frac{1}{5}(\frac{\lambda_j - \lambda_k}{h_{n,k}})^2\right]^+,$$

with $[.]^+$ denoting the positive part. In the case of $p > n$, without losing any generality, if the first n eigenvalues are zero as in equation (C.6) of Supplement of Ledoit and Wolf (2020), we have the following estimate for the spectral density

$$\hat{f}(\lambda_j) := \frac{1}{p}\sum_{k=p-n+1}^{p}\frac{3}{4\sqrt{5}h_{n,k}}\left[1 - \frac{1}{5}(\frac{\lambda_j - \lambda_k}{h_{n,k}})^2\right]^+,$$

4. Then set the diagonal matrix $\hat{\Delta}^*$

$$\hat{\Delta}^* := diag(\hat{\phi}^*(\lambda_1), \cdots, \hat{\phi}^*(\lambda_p)),$$

where if $\lambda_j = 0$ in case of $p > n$

$$\hat{\phi}^*(0) = \frac{1}{\pi(\hat{c}-1)\hat{H}_{\bar{f}}(0)}, \tag{A.1}$$

otherwise (With $\lambda_j > 0$ in case of $p < n$, or $\lambda_j > 0$ in case of $p > n$)

$$\hat{\phi}^*(\lambda_j) := \frac{\lambda_j}{[\pi\hat{c}\hat{f}(\lambda_j)]^2 + [1 - \hat{c} - \pi c\lambda_j\hat{H}_f(\lambda_j)]^2}. \tag{A.2}$$

5. Form the optimal shrinkage estimator for the covariance matrix as

$$\hat{\Sigma} := U\hat{\Delta}^*U'.$$

6. Invert $\hat{\Sigma}$ to get the precision matrix estimate $\hat{\Gamma} := (\hat{\Sigma})^{-1}$.

C Additional Material

The following subsections consider a larger time-span, a different strategy for communication between agents, short caps analysis, long-short strategies, LLM-S prompts and outputs, and a more traditional way of screening stocks and empirical results related to that.

C.1 10 Year Results

The following Tables C.1-C.8 denote the results of our model applied to test data from January 2015 to April 2024. The results are explained in Section 5.4.

BASELINE-ONLY WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.6756	0.6842	0.6900	0.0926	0.0921	0.0931	0.0188	0.0181	0.0182
Residual NW	0.3951	0.4204	0.1471	0.0492	0.0512	0.0243	0.0155	0.0148	0.0274
Deep learning	-0.3343	-0.3376	-0.2314	-65.0273	-35.9402	-0.2175	37830	11332	0.8834
POET	0.6381	0.6830	**0.7187**	0.0744	0.0779	0.0820	0.0136	0.0130	0.0130
NLS	0.6772	0.6806	0.6880	0.0798	0.0793	0.1288	0.0139	0.0136	0.0350

Table C.1: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to all firms in the S&P500. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.5955	0.5919	0.5707	0.0871	0.0843	0.0821	0.0214	0.0203	0.0207
Residual NW	0.4987	0.4680	0.2942	0.0649	0.0619	0.0496	0.0169	0.0175	0.0284
Deep learning	0.6359	0.6291	0.6269	0.0846	0.0843	0.0836	0.0177	0.0180	0.0178
POET	0.6138	0.6062	0.5510	0.0822	0.0803	0.0769	0.0180	0.0176	0.0195
NLS	**0.6367**	0.5986	0.4691	0.0820	0.0782	0.0740	0.0166	0.0171	0.0258

Table C.2: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that the LLM has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LOGISTIC REGRESSION WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.6111	0.6179	0.6604	0.0844	0.0839	0.0893	0.0191	0.0185	0.0183
Residual NW	0.4335	0.4872	0.4925	0.0537	0.0597	0.0961	0.0153	0.0150	0.0381
Deep learning	0.7244	0.7283	**0.8531**	0.0951	0.0945	0.1138	0.0172	0.0168	0.0178
POET	0.6133	0.6488	0.6397	0.0765	0.0797	0.0823	0.0156	0.0151	0.0165
NLS	0.5227	0.4938	0.6806	0.0661	0.0623	0.1443	0.0160	0.0159	0.0449

Table C.3: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that logistic regression has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

HUMAN ANALYSTS WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.5816	0.5580	**0.6181**	0.0844	0.0790	0.0875	0.0211	0.0200	0.0200
Residual NW	0.1812	0.1813	0.1860	0.0227	0.0225	0.0335	0.0157	0.0154	0.0324
Deep learning	0.5106	0.4862	0.6129	0.0668	0.0632	0.0799	0.0171	0.0169	0.0170
POET	0.4471	0.3877	0.4720	0.0598	0.0517	0.0636	0.0179	0.0178	0.0182
NLS	0.4209	0.4047	0.3787	0.0518	0.0498	0.0753	0.0151	0.0151	0.0395

Table C.4: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that analysts have screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

FINBERT WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.7406	0.7095	0.7803	0.1239	0.1643	0.1252	0.0280	0.0536	0.0258
Residual NW	0.6905	0.7279	0.8162	0.1072	0.1655	0.1242	0.0241	0.0517	0.0232
Deep learning	0.7862	0.7481	**0.8817**	0.1243	0.1675	0.1322	0.0250	0.0501	0.0225
POET	0.7444	0.7323	0.8063	0.1289	0.1710	0.1345	0.0300	0.0545	0.0278
NLS	0.7495	0.7472	0.8545	0.1131	0.1665	0.1308	0.0228	0.0496	0.0234

Table C.5: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that FinBERT has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S + HUMAN ANALYSTS WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.3969	0.1899	0.2772	0.0688	0.0375	0.0488	0.0301	0.0390	0.0309
Residual NW	0.4506	0.2283	0.1093	0.0735	0.0419	0.0224	0.0266	0.0337	0.0419
Deep learning	0.4055	0.2200	0.2036	0.0674	0.0413	0.0362	0.0276	0.0352	0.0316
POET	0.4022	0.2696	0.3127	0.0712	0.0530	0.0566	0.0314	0.0387	0.0328
NLS	**0.4953**	0.2877	0.1505	0.0833	0.0526	0.0314	0.0283	0.0335	0.0435

Table C.6: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that LLM+analysts have screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S + FINBERT WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.8262	0.2505	0.7930	0.1490	0.1174	0.1312	0.0325	0.2196	0.0274
Residual NW	0.7788	0.2146	0.6425	0.1317	0.1005	0.1118	0.0286	0.2192	0.0303
Deep learning	0.9188	0.2728	**0.9429**	0.1572	0.1275	0.1462	0.0293	0.2184	0.0241
POET	0.8338	0.2358	0.8047	0.1445	01092	0.1303	0.0300	0.2146	0.0262
NLS	0.8290	0.2282	0.6687	0.1414	0.1072	0.1198	0.0291	0.2209	0.0321

Table C.7: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that FinBERT+LLM has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

LLM-S + FINBERT + HUMAN ANALYSTS WITH QUANTITATIVE WEIGHTING: 2015-2024

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.4641	0.1022	0.3111	0.0772	0.0190	0.0525	0.0277	0.0345	0.0285
Residual NW	0.2969	-0.0284	-0.1015	0.0461	-0.0048	-0.0207	0.0241	0.0290	0.0418
Deep learning	0.4200	0.0983	0.1869	0.0671	0.0174	0.0317	0.0255	0.0315	0.0287
POET	**0.4843**	0.1685	0.3694	0.0815	0.0312	0.0640	0.0283	0.0342	0.0300
NLS	0.3679	0.0562	-0.0065	0.0596	0.0096	-0.0014	0.0263	0.0294	0.0433

Table C.8: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that LLM+FinBERT+analysts have screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

C.2 Consensus rules

As a robustness check, we modify our consensus rule for the period between January 2020-April 2024. Instead of pivoting to the union of LLM-S and FinBERT when their intersection is null, we go with FinBERT's recommendation, since it achieves a higher Sharpe ratio by itself then LLM-S in most cases in Tables 2 and 5. Table C.9 presents the results. Although the winner does not achieve as high of a SR as the main table in Table 7 (with union), there are six method-portfolio combinations that do better. Overall, the magnitude of the Sharpe ratios are relatively close to the default union, showing robustness.

	Sharpe Ratio			**Returns**			**Variance**		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.9262	1.0596	0.8688	0.2299	0.3416	0.1939	0.0616	0.1039	0.0498
Residual NW	0.9593	1.0223	0.8804	0.2174	0.3237	0.1771	0.0513	0.1003	0.0405
Deep learning	0.9917	1.0478	0.9827	0.2359	0.3335	0.2017	0.0566	0.1013	0.0421
POET	0.9287	**1.0758**	0.8989	0.2356	0.3380	0.2103	0.0644	0.0987	0.0547
NLS	0.9683	1.0074	0.9477	0.2245	0.3261	0.1884	0.0537	0.1048	0.0395

Table C.9: Annualized Sharpe ratios, returns, and variance with different methods of estimating the precision matrix, with different objective functions, applied to firms that FinBERT+LLM has screened. Here, the consensus rule is to use the intersection, and defaulting to FinBERT's recommendation in case of a null intersection. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

C.3 Small Caps

We have also considered the question of how small caps factor into our strategy and Sharpe ratio between January 2020-April 2024 out-sample period. We have also analyzed the bottom 100 stocks by market capitalization in our universe, and applied our buy-sell decisions corresponding to them only in our main table, which contains LLM-S, FinBERT and the quantitative estimation method with the Agentic AI structure. The Sharpe ratios decrease. For example, nodewise/GMV decreases from 0.9212 in Table 7 to 0.6553, deep learning/GMV decreases from 1.0148 to 0.6430, deep learning/MV decreases from 1.1867 to 0.7920, and deep learning/MSR decreases from 1.0107 to 0.8703. Similar decreases also happen for NLS.

Our explanation for this behavior is that since we are now restricting to firms with the 100 smallest market capitalizations in the S&P500, this contradicts with most buy signals from LLM-S (since most buy signals invest in firms with high market caps, as illustrated in the example in Section 2.1), which may explain why the Sharpe ratio decreases.

C.4 Long Short Strategies

In this section, we investigate equal- and value-weighted long-short strategies using our models. There are two possible strategies used, first LLM-S based long-short portfolio, and then Best-2

stage based long-short portfolio which will be explained below. For each strategy we analyze both 5 and 10 year periods that we analyzed.

First, we evaluate the long-short performance on the LLM-S's buy and sell recommendations. We construct a market-neutral equal-weighted portfolio invested in all the buy and sell recommendations. Similarly, we can use a size-based value-weighted portfolio as well. We do this for both 5 year and 10 year windows.

A strategy based on LLM-S agent screened stocks on equally weighted long-short portfolio-provides in Table C.10, a Sharpe Ratio of -0.7805 in the five year period that we analyzed, namely between January 2020-April 2024. That SR is much lower than 0.6738 in Table 2 where LLM-S agent results are used in a nonlinear shrinkage GMV portfolio. We also consider a value-weighted portfolio based on firm size, which achieves a SR of -1.0463, also much lower than 0.6738.

The second strategy is best-2 stage model, where instead of only LLM-S we use the Finbert+LLM-S+quantitative method to rank the weights from the largest to the smallest one. Then, since the quantitative strategy weights are numerical scores, we can rank them from the largest to smallest. Specifically, we take the top and bottom deciles of scores to long and short, respectively. We also consider equal and value weighted portfolios, where the latter is proportional to weights that the Finbert+LLM-S+quantitative method outputs.

The best 2-stage based long-short portfolio uses MV portfolio in case of a 5 year window, and MSR portfolio for a 10 year window. For equally weighted 10 year window, we obtain 0.1508 as SR, compared to SR of 0.9429 via Finbert plus LLM-S screened and then used Deep learning method to get portfolio weights in Table C.7. We also consider a value-weighted portfolio based off of the quantitative method weights. For this portfolio, we also obtain a lower SR.

It is clear from our Table C.10, that a simple long-short equally weighted or value weighted strategy delivers low SR compared with what we propose. So this also shows the benefit of using Agentic-AI based data dependent screening and weighting makes a difference compared to ranking of all stocks and then forming a portfolio based on top-bottom-deciles.

	5 year		**10 year**	
Model	**Equal-Weighted**	**Value-Weighted**	**Equal-Weighted**	**Value-Weighted**
LLM-S	-0.7805	-1.0463	-0.5581	-0.7685
Best 2 stage model	-0.0834	0.0612	0.1508	0.0475

Table C.10: Annualized Sharpe ratios incorporating long-short strategies for both 5 and 10-year evaluation windows. The best 2 stage model is Finbert+LLM-S+quantitative method using MV for the 5 year window, and the same agent using MSR for the 10 year window.

C.5 LLM-S Prompts and Outputs

Below is a code snippet describing the LLM-S agent. This snippet is specifically for test dates in 2024.

```
strategy_agent = Agent(
    role="Quantitative Strategy Developer",
    goal="Develop systematic BUY/HOLD/SELL rules based on firm characteristics that can be applied
    ↪  to all S&P 500 firms",
    backstory="""
    You are an expert quantitative strategist who creates systematic, rule-based trading
    ↪  strategies.

    CRITICAL DATA UNDERSTANDING:
    - 'mve' = log(market value of equity), represents log firm size
    - 'bm' = book-to-market ratio (value factor). Understand that a high book-to-market value means
    ↪  undervalued, and a low book-to-market value means overvalued.
    - 'mom12m' = 12-month momentum
    - ALL features are standardized: mean = 0, standard deviation = 1
    - Values are z-scores showing standard deviations from mean

    Your task is to develop EXPLICIT, SYSTEMATIC RULES for generating trading signals. Understand
    ↪  that when doing the following, you must use causal masking from December 2023 to prevent
    ↪  any look-ahead bias.

    1. EXPLORE THE DATA (December 2023):
       - Identify what constitutes "extreme" values for mve, bm, and mom12m
       - Look for natural clustering or breakpoints in the data
       - Consider correlations between characteristics

    2. DEVELOP CLEAR RULES based on economic intuition:
       - Keep in mind the market conditions at this date - this might influence the rules you
       ↪  choose.
       - The following are example questions you can consider, BUT THEY ARE NOT EXHAUSTIVE:
           - Value stocks: Should low bm (cheap) be BUY or SELL?
           - Momentum: Should high mom12m (strong performance) be BUY or SELL?
           - Size: Should mve matter for the strategy?
           - Combinations: What about value + momentum + size together?

    3. DEFINE SPECIFIC THRESHOLDS:
     Your output must include exact rules. You can use test_complex_condition to test different
     ↪  combinations.
       Your output must include exact rules like:
       - "BUY if: bm < -0.71 AND mom12m > 0.85 AND mve > 0.23" or "BUY if: (bm > 0.57 AND mom12m <
       ↪  0.82) OR mve > -0.88" or "BUY if: bm > 0.63 OR mve < 0.98"
       - "SELL if: bm > 1.25 OR mom12m < -0.98" or "SELL if: (bm < -0.84 OR mve < 0) AND mom12m <
       ↪  -0.97" or "SELL if: (bm > 0.94 AND mom12m < -1.06) OR mve <-0.68"
```

```
   - "HOLD: all other cases"
   - However, the above is only AN EXAMPLE - so do not simply copy the format above. You are
   ↪  free to include/exclude as many conditions in the if statements. You are also free to
   ↪  make the conditions as complicated or as simple as you like.
   - Be PRECISE in your thresholds, do not choose numbers that are nice or round - you are a
   ↪  quantative strategy developer.

4. PROVIDE RATIONALE:
   - Why these thresholds?
   - What's the economic intuition?
   - What patterns did you observe in the data?

CRITICAL REQUIREMENTS:
- Rules must be DETERMINISTIC (same inputs → same output)
- Use ONLY z-score comparisons (>, <, AND, OR), but these may be impacted by market conditions.
- Define thresholds for BUY, SELL, and HOLD
- Rules should be implementable as: if (condition) then signal
- Keep in mind that we will use the buy signals to construct a portfolio, so it is better to
↪  give too many signals, rather than too few signals.

OUTPUT FORMAT:
==========================================
SYSTEMATIC TRADING RULES
==========================================

Data Exploration Summary:
- [Key statistics and patterns observed]

BUY RULE:
if [ANY complex z-score condition using AND/OR/NOT]:
    signal = BUY

Examples of valid BUY rules:
- "bm < -1.15 AND mom12m > 0.73 AND mve < 1.11" (simple AND)
- "bm < -1.56 OR mom12m > 1.28 OR mve > 1.52" (simple OR)
- "(bm < -1.02 AND mom12m > 0.53) OR mve > 1.59" (combination)
- "bm < -0.83 AND (mom12m > 0.77 OR mve > 1.08)" (nested conditions)

SELL RULE:
if [ANY complex z-score condition]:
    signal = SELL

HOLD RULE:
else:
    signal = HOLD

Rationale:
```

```
    - [Economic reasoning for BUY rule]
    - [Economic reasoning for SELL rule]
    - [Expected signal distribution]

    ```
    ==========================================

    Be precise, systematic, and data-driven. Your rules will be applied to ~500 firms.
    """
)
```

Further, below is code snippet on the task description of our LLM-S model.

```
strategy_task = Task(
    description="""
    Develop systematic BUY/HOLD/SELL rules for S&P 500 firms at December 2023.

    Available characteristics (all are z-scores):
    - mve: log firm size
    - bm: log book-to-market (value)
    - mom12m: 12-month momentum

    Your process:
    1. Get database schema and understand available data for December 2023
    2. Explore extreme values for each characteristic
    3. Look for patterns and correlations
    4. Use test_complex_condition to test different rule combinations
    5. Develop systematic rules with specific z-score thresholds

    You have COMPLETE FLEXIBILITY in creating rules. Test different combinations using AND, OR,
    ↪  NOT.

    CRITICAL: Your output must be EXPLICIT RULES with exact thresholds that can be
    implemented in Python/pandas. You must give PRECISE THRESHOLDS - do not give thresholds that
    ↪  are only nice or round numbers. Further, use causal masking from December 2023 to prevent
    ↪  any look-ahead bias.

    Focus on:
    - Economic intuition (value, momentum, size effects)
    - Clear, implementable thresholds
    - Balance between signal strength and diversification
    - Rules that make sense for ~500 firms

    Output systematic rules that I can directly implement in code.
    """,
    expected_output="""
```
```

```
    Complete strategy document with:
    1. Data exploration summary
    2. Explicit BUY rule with z-score thresholds
    3. Explicit SELL rule with z-score thresholds
    4. HOLD rule (default case)
    5. Economic rationale for each rule

    Rules must be deterministic and implementable.
    """,
    agent=strategy_agent,
    tools=[
        get_database_schema,
        query_firm_database,
        get_extreme_firms,
        test_complex_condition
    ]
)
```

Lastly, below is an example of the corresponding output of the LLM-S model. This particular year, the LLM likes buying undervalued firms with above-average size and non-small momentum values. Conversely, it likes selling firms that are overvalued, small firms, or negative-momentum firms (but if they have not shown significantly positive momentum).

```
==========================================
SYSTEMATIC TRADING RULES
==========================================

Data Exploration Summary:
- The data includes standardized book-to-market (bm), momentum (mom12m), and market value of equity
↪  (mve) for S&P 500 firms as of December 2023.
- All features are standardized to have a mean of 0 and a standard deviation of 1.
- Extreme values were explored to understand the distribution of characteristics.

BUY RULE:
if bm > 0.95 AND mve > 0.3 AND mom12m > -0.5:
    signal = BUY

SELL RULE:
if (bm < -0.75 OR mom12m < -0.55 OR mve < -0.75) AND NOT (mom12m > 1.5):
    signal = SELL

HOLD RULE:
else:
    signal = HOLD

Rationale:
```

```
- BUY Rule: This rule targets undervalued (high bm), reasonably sized (mve > 0.3) companies with
↪  positive momentum (mom12m > -0.5). The economic intuition is to buy companies that are
↪  currently cheap but have shown some signs of recovery or positive market sentiment.
- SELL Rule: This rule aims to sell companies that are overvalued (low bm), have negative momentum
↪  (mom12m < -0.55), or are small in size (mve < -0.75). The `NOT (mom12m > 1.5)` condition
↪  prevents selling companies with extremely high momentum, even if they meet the other criteria,
↪  as these might be temporary situations or represent significant growth opportunities. The
↪  economic intuition is to avoid holding onto companies that are losing value or are too small to
↪  provide substantial returns, unless they are exhibiting exceptional positive momentum.
- Expected signal distribution: The BUY rule is expected to generate signals for a small percentage
↪  of firms (around 2%), focusing on higher-conviction value opportunities. The SELL rule is
↪  expected to affect a larger percentage of firms (around 50%), filtering out less desirable
↪  investments and managing risk. The remaining firms will be held.
```

C.6 Novy-Marx Screening

In this part, we analyze Novy-Marx (2013) based screening effect on Sharpe Ratio analysis. Table C.11 contains Sharpe ratios, returns, and variances with the screening method described in Novy-Marx (2013), with out-of-sample test periods from January 2020 to April 2024. Table C.12 uses out-of-sample test periods from January 2015 to April 2024. Out of 500 firms, we take the 150 stocks with the highest combined profitability and value ranks, and the 150 stocks with the lowest, to be the screened set of stocks each year. The quantitative method then assigns weights to form a portfolio, whose Sharpe ratios we report in the tables.

To see the effect that LLM-S has in the screening ensemble containing LLM-S and FinBERT, we replace LLM-S instead with the method described in Novy-Marx (2013). Notice that in the 5 year sample in Table C.13, the Sharpe ratio is lower than the one obtained by the LLM-S plus FinBERT ensemble when compared with Table 7. The same result holds true for the 10 year sample in Table C.14 comparison with Table C.7 as well. This suggests another contribution of our LLM-S screening model: that it has better synergy with FinBERT than with other screening methods.

Further, note that single-agent LLM-S screening itself beats single-agent Novy-Marx screening in the 5 year window, but not the 10 year window. This is somewhat expected, since Novy-Marx picks out growth stocks, which are suited for the 10 year window. Regardless, the ensemble containing both LLM-S and FinBERT beat Novy-Marx in both windows, showing that FinBERT can also decrease mistakes made by LLM-S when acting as an ensemble.

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.3659	0.3743	0.3901	0.0627	0.0627	0.0646	0.0294	0.0281	0.0275
Residual NW	0.0250	0.0416	0.0833	0.0036	0.0058	0.0176	0.0203	0.0192	0.0447
Deep learning	0.4461	0.4354	0.4500	0.0797	0.0760	0.0777	0.0319	0.0305	0.0298
POET	0.1887	0.1963	0.1985	0.0283	0.0290	0.0302	0.0224	0.0218	0.0232
NLS	0.3151	0.3398	**0.5069**	0.0430	0.0453	0.1150	0.0187	0.0178	0.0515

Table C.11: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that Novy-Marx has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.5931	0.6004	0.6190	0.0839	0.0833	0.0851	0.0200	0.0192	0.0189
Residual NW	0.3745	0.4210	0.3554	0.0471	0.0518	0.0600	0.0158	0.0151	0.0285
Deep learning	0.6652	0.6692	**0.6904**	0.0960	0.0944	0.0966	0.0208	0.0199	0.0196
POET	0.5665	0.5916	0.6048	0.0705	0.0724	0.0754	0.0155	0.0150	0.0156
NLS	0.5309	0.5512	0.4442	0.0633	0.0645	0.0829	0.0142	0.0137	0.0348

Table C.12: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that Novy-Marx has screened. Test period is from Jan. 2015-April 2024. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.2554	0.5570	0.1519	0.0558	0.1442	0.0313	0.0478	0.0670	0.0425
Residual NW	0.2741	0.6695	0.1955	0.0556	0.1772	0.0427	0.0412	0.0701	0.0478
Deep learning	0.3340	0.6485	0.2335	0.0701	0.1635	0.0473	0.0440	0.0636	0.0411
POET	0.2399	0.5238	0.1841	0.0539	0.1322	0.0389	0.0505	0.0637	0.0446
NLS	0.3206	**0.6929**	0.2208	0.0642	0.1773	0.0490	0.0401	0.0655	0.0493

Table C.13: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that Novy-Marx + FinBERT has screened. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

	Sharpe Ratio			Returns			Variance		
Method	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**	**GMV**	**MV**	**MSR**
NW	0.3815	0.6054	0.3332	0.0678	0.1272	0.0564	0.0316	0.0442	0.0287
Residual NW	0.3362	0.6286	0.2813	0.0565	0.1335	0.0504	0.0282	0.0451	0.0321
Deep learning	0.3980	**0.6385**	0.3524	0.0682	0.1311	0.0587	0.0294	0.0422	0.0278
POET	0.3361	0.5555	0.3043	0.0608	0.1148	0.0523	0.0327	0.0427	0.0295
NLS	0.3646	0.6345	0.2625	0.0605	0.1314	0.0480	0.0276	0.0429	0.0334

Table C.14: Annualized Sharpe ratios with different methods of estimating the precision matrix, with different objective functions, applied to firms that Novy-Marx + FinBERT has screened. Test period is from Jan. 2015-April 2024. GMV=Global minimum variance portfolio, MV=mean-variance portfolio with target returns as 1% monthly, MSR=maximum Sharpe ratio portfolio.

References

Altman, E. (1968). Financial ratios, discriminant analysis and the predictions of corporate bankruptcy. *Journal of Finance 23*, 589–609.

Araci, D. (2019). Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*.

Arvanitis, S., O. Scaillet, and N. Topaloglou (2024). Sparse spanning portfolios and under-diversification with second-order stochastic dominance. Technical Report 2402.1951, arxiv.

Asness, C. S., A. Frazzini, and L. H. Pedersen (2019). Quality minus junk. *Review of Accounting studies 24*(1), 34–112.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*(1), 191–221.

Bishop, C. and H. Bishop (2024). *Deep Learning, Foundation and Concepts*. Springer.

Callot, L., M. Caner, O. Onder, and E. Ulasan (2021). A nodewise regression approach to estimating large portfolios. *Journal of Business and Economic Statistics 39*, 520–531.

Caner, M. and M. Daniele (2025). Deep learning based residuals in non-linear factor models: Precision matrix estimation of returns with low signal-to-noise ratio. *Journal of Econometrics 251*, 106083.

Caner, M. and Q. Fan (2025). A practioner's guide to ai-ml in portfolio investing. Technical Report 2509.25456v1, arxiv.

Caner, M. and Q. Fan (2026). Portfolio analysis in high dimensions with tracking error and weight constraints. *Journal of The American Statistical Association, Forthcoming*.

Caner, M., M. Medeiros, and G. Vasconcelos (2023). Sharpe ratio analysis in high dimensions: Residual based nodewise regression in factor models. *Journal of Econometrics 235*, 393–417.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of finance 52*(1), 57–82.

Chang, J., Y. Qiu, Q. Yao, and T. Zou (2018). Confidence regions for entries of a large precision matrix. *Journal of Econometrics 206*(1), 57–82.

Chen, A. Y. and J. McCoy (2024). Missing values handling for machine learning portfolios. *Journal of Financial Economics 155*, 103815.

Chen, Y., B. Kelly, and D. Xiu (2023). Expected returns and large language models. Technical Report https://ssrn.com/abstract=4416687, SSRN.

Didisheim, A., M. Fraschini, and L. Somoza (2025). Ai's predictable memory in financial analysis. *Economics Letters 256*, 112602.

Fama, E. and K. French (1992). The cross-section of expected stock returns. *Journal of Finance 47*, 427–465.

Fan, J., Y. Liao, and M. Mincheva (2011). High-dimensional covariance matrix estimation in approximate factor models. *The Annals of Statistics 39*, 3320–3356.

Fan, J., Y. Liao, and M. Mincheva (2013). Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 75*(4), 603–680.

Fan, Y. and C. Y. Tang (2013). Tuning parameter selection in high dimensional penalized likelihood. *Journal of the Royal Statistical Society Series B: Statistical Methodology 75*(3), 531–552.

Farrell, M. H., T. Liang, and S. Misra (2021). Deep neural networks for estimation and inference. *Econometrica 89*(1), 181–213.

Green, J., J. R. Hand, and X. F. Zhang (2017). The characteristics that provide independent information about average us monthly stock returns. *The Review of Financial Studies 30*(12), 4389–4436.

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

Hendrycks, D., C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt (2020). Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.

Hong, H., T. Lim, and J. Stein (2000). Bad news travels slowly: size, analyst coverage, and the profitability of momentum strategies. *Journal of Finance 55*, 265–295.

Horn, R. and C. Johnson (2013). *Matrix Analysis*. Cambridge University Press.

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

Kamath, U., K. Keenan, G. Somers, and S. Sorenson (2024). *Large Language Models: A Deep Dive*. Springer.

Kelly, B. and D. Xiu (2023). Financial machine learning. *Foundations and trends in finance 13*, 205–363.

Kelly, B. T., B. Kuznetsov, S. Malamud, and T. A. Xu (2025). Artificial intelligence asset pricing models. Technical report, National Bureau of Economic Research.

Kozak, S., S. Nagel, and S. Santosh (2020). Shrinking the cross-section. *Journal of Financial Economics 135*(2), 271–292.

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

Ledoit, O. and M. Wolf (2020). Analytical nonlinear shrinkage of large-dimensional covariance matrices. *The Annals of Statistics 48*(5), 3043–3065.

Li, W., H. Kim, M. Cucuringu, and T.Ma (2025). Can llm-based financial investing strategies outperform the market in lomg-run. Technical Report 2505.07078v2, arxiv.

Li, Y., S. Wangand, H. Ding, and H. Chen (2024). Large language models in finance: A survey. Technical Report 2311.10723v2, arxiv.

Ludwig, J., S. Mullainathan, and A. Rambachan (2025). Large language models: An applied econometric framework. Technical Report 2402.07031v3, arxiv.

Malo, P., A. Sinha, P. Korhonen, J. Wallenius, and P. Takala (2014). Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the Association for Information Science and Technology 65*(4), 782–796.

Meinshausen, N. and P. Bühlmann (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, 1436–1462.

Mohanram, P. S. (2005). Separating winners from losers among lowbook-to-market stocks using financial statement analysis. *Review of accounting studies 10*(2), 133–170.

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

Piotroski, J. (2000). Value investing: The use of historical financial statement information to separate winners from losers. *Journal of Accounting Research 38*, 1–41.

Sapkota, R., K. Roumeliotis, and M. Karkee (2026). Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges. *Information Fusion 126*, 103599.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: heuristics and bias. *Science 185*, 1124–1131.

Tversky, A. and D. Kahneman (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review 90*, 293–315.

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Poloshukhin (2017). Attention is all you need. Technical Report 1706.03762, arXiv.

Zhang, Q., C. Hu, S. Upasani, B. Ma, F. Hong, V. Kamanuru, J. Ranton, C. Wu, M. Ji, H. L. U. Thakker, J. Zou, and K. Olukotun (2025). Agentic context engineering:evolving contexts for self improving language models. Technical Report 2510.04618v1, arxiv.

Zhao, T., J. Lyu, S. Jones, H. Garber, S. Pasquali, and D. Mehta (2025). Alpha agents: Large language model based multi-agents for equity portfolio constructions. Technical Report 2508.11152v1, arxiv.