Learning from the Wisdom of Mutual Fund Managers*

Roméo Tédongap[†] Jules Tinang[‡]

June 4, 2025

Abstract

We define Stock Active Share (SAS) as the degree to which a stock in a benchmark index is actively weighted by mutual funds relative to its index weight. We analyze the risk-return characteristics of portfolios ranked by SAS values. The top quantile portfolio delivers significant monthly risk-adjusted returns, highlighting mutual fund managers' capital allocation proficiency. However, due to the delayed disclosure of fund holdings, SAS is unobservable in real-time, making the strategy unfeasible for typical investors. To address this, we apply machine learning models to historical fund holdings and stock characteristics to predict future SAS and sort portfolios accordingly. These models demonstrate substantial out-of-sample accuracy, and the feasible top quantile portfolio consistently outperforms the benchmark across risk-adjusted measures. Our findings illustrate the enduring value of fund managers' stock-picking skills, challenging the view that technological advancements diminish their importance. Furthermore, the feasible strategy outperforms traditional analyst recommendations and aligns with sustainability goals by favoring stocks with lower carbon intensity.

Keywords: Asset Pricing, Machine Learning, Firm Characteristics, Mutual Funds.

JEL Classification: G11, G17, G23.

*We thank Javier Gil-Bazo, Lammertjan Dam, Auke Plantinga, participants at the 7th Asset Pricing Conference by LTI@UniTO at Collegio Carlo Alberto, Italy, and the ESSEC-Amundi Webinar "On the Use of AI and Machine Learning in Asset Management."

[†]Corresponding author: Professor of Finance, Department of Finance, ESSEC Business School, 3 Avenue Bernard Hirsch, CS 50105 CERGY, 95021 Cergy-Pontoise Cedex, France. Email: tedongap@essec.edu.

[‡]Visiting Independent Researcher, Department of Finance, ESSEC Business School, 3 Avenue Bernard Hirsch, CS 50105 CERGY, 95021 Cergy-Pontoise Cedex, France. Email: jules.tinang@essec.edu.

1 Introduction

The digital transformation of modern economies has dramatically improved access to information and financial markets. Ordinary investors now have easy access to trading platforms, often through smartphones, allowing them to manage their portfolios (Eaton et al.; 2022; Bryzgalova et al.; 2023). This increased accessibility, however, raises questions about the role of professional investors who manage assets on behalf of others. Specifically, it remains unclear whether mutual funds have a reliable source of information that allows them to allocate capital effectively to firms capable of generating positive risk-adjusted returns. Additionally, a practical question arises: can investors learn and profit from the skills of mutual fund managers by analyzing the history of mutual fund holdings? This paper provides evidence supporting a positive answer to this question.

We consider Stock Active Share (SAS) to measure the degree to which a stock belonging to a standard benchmark index is actively weighted by mutual funds compared to its weight in the benchmark index, and we analyze the risk-return characteristics of portfolios formed by ranking stocks according to their active shares among mutual funds. SAS for a given stock in a benchmark index, calculated as the sum of the absolute differences between the weights of the stock in various mutual funds and its weight in the benchmark index, is interpreted as a measure of managers' conviction about the stock's value. We focus on two benchmark indexes: the MidCap S&P400 and the BigCap S&P500. We show that knowing ex-ante the information that mutual funds reveal ex-post about their portfolio composition, like an Oracle, allows us to design a SAS-based investment strategy that consistently outperforms the market index. This strategy, termed our "SAS-Oracle strategy," which uses quarterly rebalancing to sort stocks into portfolios based on SAS values and buys the top quantile portfolio, generates monthly risk-adjusted returns ranging across benchmark indexes from 0.67% to 0.92% in tercile portfolios, from 0.73% to 1.15% in quintile portfolios, and from 0.74% to 1.41% in decile portfolios.

Although the findings suggest that mutual fund managers can efficiently allocate capital and surpass passive investors, their skills do not directly benefit mutual funds mimicking investors due to the delayed release of information on mutual fund holdings. When this information is available, it's too late for mutual funds mimicking investors to use it to implement the SAS-Oracle strategy effectively. We employ machine learning models to anticipate SAS value to address this limitation. We train four machine learning models—elastic net, random forest, gradient boosting, and deep neural network—and utilize ordinary least squares (OLS) regression on historical data encompassing thousands of month-stock observations and hundreds of stock characteristics, aiming to predict mutual fund managers' later-released conviction about stock value. This predictive approach allows us to construct our investment portfolios proactively.

We initially investigate the capacity of machine learning models to predict the SAS value of stocks, our primary measure of interest. Our analysis reveals a notably high out-of-sample Pearson correlation between the actual SAS values and their predictions by the machine learning models. For instance, we see a 66% correlation for deep neural network (DNN) and 70% for random forest (RF) in BigCap stocks, and 46% for DNN and 44% for RF in MidCap stocks, with the models being updated every three months. Similar levels of accuracy are seen in other goodness-of-fit measures, such as out-of-sample R-squared values. These metrics are crucial, as higher out-of-sample fits indicate that sorting stocks based on their SAS predictions by the machine learning models can effectively proxy portfolio sorts based on actual but unobservable SAS values.

Next, we demonstrate the efficacy of a machine learning-based investment strategy that uses quarterly rebalancing to sort stocks into portfolios based on predicted SAS values, explicitly targeting the top quantile portfolio. This strategy termed our "SAS-feasible strategy," consistently outperforms the benchmark index across various measures of risk-adjusted returns. For BigCap stocks, the monthly risk-adjusted returns achieved by the DNN and RF models in the SAS-feasible top quintile portfolio are 0.35% and 0.30%, respectively. These figures represent 86.93% and 74.37% of the risk-adjusted return in the SAS-Oracle top quintile portfolio during the same period. In the MidCap category, the DNN and RF models deliver monthly risk-adjusted returns of 0.62% and 0.51%, respectively, corresponding to 84.27% and 69.90% of the SAS-Oracle top quintile portfolio counterparts. For comparison, the two benchmark indexes record a zero monthly risk-adjusted return during the same period.

Our strategy allows us to create relatively concentrated yet diversified portfolios anytime. We observe that the total risk of the SAS-feasible strategy is comparable to, if not lower than, that of the benchmark index while offering higher returns. Additionally, our portfolio strategy exhibits less negative skewness and lower kurtosis than the benchmark indexes, indicating a lower probability of adverse outcomes for active investors than those passively holding the benchmark index. Our findings are robust across different quantile-based portfolio formations (tercile, quintile, or decile), allowing us to assess the impact on diversification by varying the number of stocks in our portfolio. Surprisingly, the portfolio performance remains very stable, making it suitable for modest individual investors and more prominent institutional investors seeking to hold more stocks.

This novel finding suggests that investors can enhance their portfolios by learning from mutual fund managers' historical holdings, which reflect skilled decision-making. A fund's portfolio results from an optimization process where a manager uses various stock characteristics and market signals to allocate capital effectively. Skilled managers consistently select valuable stocks, deviating from the market index to achieve higher returns. If mutual fund performance were purely driven by luck, a deterministic strategy based on fund holdings wouldn't consistently outperform the market. This paper demonstrates that fund managers' collective skills provide valuable information, enabling profitable asset allocation strategies that generate positive risk-adjusted returns for investors.

This article connects to several strands of literature. The debate on mutual funds' ability to select well-performing stocks and time the market remains unsettled. Early research found little evidence of manager skill, leading to skepticism about the value of active management and a preference for passive investing (Carhart; 1997; Fama and French; 2010). The conventional view suggests that most funds underperform after fees, with limited persistence in performance and few managers showing skill beyond costs (Cremers and Petajisto; 2009). However, recent studies using different measures, such as active shares (Daniel et al.; 1997; Cremers and Petajisto; 2009) and funds' inflows (Berk and van Binsbergen; 2015), challenge this view, finding that some mutual funds can outperform the market even after fees, and that top-performing funds often exhibit persistence.

Our paper builds on Kacperczyk et al. (2005) and Cremers and Petajisto (2009) who predict mutual fund performance through their aggregate deviation from the fund-specific benchmark weights. Jones and Mo (2020) investigate the out-of-sample performance of mutual fund predictors and discover that related publications rendered mutual fund active share and industry concentration ineffective in predicting mutual fund performance. They explain this through a learning channel, meaning investors or mutual fund managers learn from academic literature and change their investments. We differ from them in two ways: first, we use stock deviations relative to standard benchmark indices rather than fund-specific benchmarks; second, we aggregate deviations across funds for each stock rather than across stocks in a fund. This approach departs from the original active share definition, and we argue it provides novel market information. Empirical results support this, as portfolios based on our SAS measure consistently outperform benchmark indices.

Our approach follows Jiang et al. (2014), who measure mutual funds' deviations from benchmarks by averaging the simple (not absolute) differences between a stock's weight in mutual funds and its weight in the benchmark index. Unlike their focus on stocks held by mutual funds or in the fund's specific benchmark, we concentrate on stocks within a standard benchmark. Like them, we calculate the Net Stock Active Share (NSAS) as the sum of the simple differences between a stock's weight in mutual funds and its benchmark. Our findings reveal that positive differences are significantly larger than negative ones, and strong positive correlations between simple and absolute differences: 0.89 for the S&P500 and 0.90 for the S&P400. Thus, both NSAS and SAS should yield similar results for stock selection.

Jiang et al. (2014) show that decile portfolios based on average deviations generate alpha, suggesting mutual fund managers possess stock-picking skills. We confirm this by demonstrating that our SAS-Oracle strategy, which selects top quantile stocks by SAS values, outperforms the market index. Unlike previous studies, we address the impracticality of delayed mutual fund holdings' disclosure by using machine-learning methods, historical fund holdings, stock characteristics, and macroeconomic data to build tradable portfolios that consistently outperform market indices.

Additionally[1], Antón et al. (2008, 2021) developed a "best ideas" portfolio for mutual fund

[1] Wermers et al. (2012), Agarwal et al. (2013), and Yan and Zhang (2007) find positive return predictability from the portfolio holdings of actively managed mutual funds, hedge funds, and short-term institutional investors, respectively.

managers, but our approach differs. Their method focuses on maximizing the Sharpe ratio for individual managers' top ideas, while our strategy seeks to identify the best ideas across managers, resulting in a more concentrated portfolio. Despite this, our approach remains flexible, allowing for adjustments in the number of stocks included. Our portfolio demonstrates high active share with low tracking error, consistent with diversified stock selection as per Cremers and Petajisto (2009).

The rise of artificial intelligence in finance has spurred numerous studies on machine learning applications in investment (Kelly and Xiu; 2023). For example, Gu et al. (2020) employ machine learning algorithms with multiple predictors to construct stock portfolios that outperform the market. Similarly, we utilize machine learning and stock characteristics to develop a profitable investment strategy. However, we further enhance this approach by incorporating insights from fund manager expertise to achieve this objective. Li and Rossi (2021) apply machine learning models and mutual fund holdings' characteristics to predict high-performing mutual funds. Kaniel et al. (2023) integrate deep learning with mutual fund characteristics, holdings, and investor sentiment to forecast mutual fund performance, concluding that stock characteristics are not essential for selecting top-performing funds. DeMiguel et al. (2023) supported this view, showing that machine learning and mutual fund characteristics can be used to construct profitable mutual fund portfolios.

In contrast, our objective is not to identify the best-performing mutual funds but to predict the best-performing stocks within benchmark indices using the aggregate information from mutual funds. Our contribution lies in leveraging machine learning and stock characteristics to predict mutual fund managers' collective stock valuations based on their deviations from a standard benchmark. Unlike traditional approaches focusing on stock returns, we use mutual funds' aggregate deviations as the target variable in our models. This allows us to capture return maximization and broader objectives like managing drawdown or value-at-risk, offering more profound insights into mutual funds' investment strategies.

We extend our performance comparison beyond standard stock indices, including comparing SAS strategies with those based on analyst recommendations. Barber et al. (2001) show that buying stocks with favorable consensus recommendations is profitable. Our results indicate a weak but positive correlation between analyst recommendations and SAS measures. While both strategies outperform benchmark indices, the SAS-based strategies consistently deliver better risk-adjusted returns, suggesting that machine learning predictions based on mutual fund manager behavior provide more effective investment guidance than analyst recommendations.[2]

Finally, we explore the relationship between fund manager stock preferences as reflected in SAS measures and carbon emissions level and intensity. On the one hand, our findings indicate a significant negative correlation between a stock's carbon emissions intensity and its SAS value, meaning that stocks favored by fund managers for outperforming benchmark indices tend to have lower carbon emissions intensity or, equivalently, to be more carbon efficient. This aligns the SAS strategy with ESG criteria, demonstrating that achieving both strong financial returns and sustainable investment goals is possible. On the other hand, we find a positive relationship between carbon emissions level and SAS value, which shows that a stock's higher level of carbon emissions predicts a higher expected return for investors, suggesting that higher SAS stocks may earn a carbon premium for their exposure to climate-related risks (Bolton and Kacperczyk; 2021, 2023).

The remainder of the paper is organized as follows: Section 2 introduces the stock active share, our measure of managers' conviction regarding stock value, and provides details on our data. Section 3 explores the composition and performance of the SAS-Oracle portfolios. Section 4 focuses on the machine learning implementation and performance for predicting SAS measures. Section 5 presents and evaluates the main results, focusing on the financial performance of the machine learning-based investment strategy. Section 6 explores the relationship between the SAS-feasible strategy and carbon emissions. Section 7 concludes the paper. An External Appendix offers additional results on portfolio rebalancing at a monthly frequency.

[2] We show in the internet appendix that sorting stocks on lagged one-month, two-month, or three-month SAS values, which are readily available, renders the investment strategy ineffective. More specifically, the performance achieved is not better than analyst recommendations and worse than investment strategies based on machine learning.

2 Data, Measures, and Descriptive Statistics

This section explores publicly available data on U.S. mutual fund holdings. It introduces and formally defines the stock active share, a crucial measure upon which all subsequent empirical analyses are based. Additionally, relevant summary statistics for stock active share are provided and discussed, serving as a foundation for the empirical investigation that follows.

2.1 Data on Mutual Fund Holdings of Stocks

We collect data from multiple sources, primarily the CRSP Survivor-Bias-Free U.S. Mutual Fund Database, covering the period from August 2007 to December 2023.[3] We begin by selecting actively managed U.S. equity mutual funds using the Fund Summary table, following the methodology of Kacperczyk et al. (2006). Funds are chosen based on their Investment Objective Codes.[4] Index funds are identified and excluded by their names,[5] and only funds with at least two-thirds of their total net assets (TNA) in common stocks are retained.

We obtain the CRSP holdings file, which provides a list of equities held by open-ended funds, along with their percentage of total net assets. The historical list of benchmark index constituents is sourced from Refinitiv/Datastream, along with the market capitalization of each constituent, enabling us to calculate the portfolio weights of the benchmark index. The weight of each stock in the benchmark index is determined by dividing its market value by the total market value of the index. These weights are then used to calculate the deviation of each stock, in the benchmark index, from the corresponding holdings reported by each mutual fund. These deviations are subsequently

[3]The SEC's Rule IC-26372, implemented in May 2004, requires mutual funds to disclose their end-of-quarter holdings four times a year within 60 days. Prior to this, disclosure was optional and limited to a few funds. Before August 2007, our dataset includes a small number of funds (fewer than 10) that disclosed holdings regularly.

[4]We select funds with the following Lipper classification codes: EIEI, G, LCCE, LCGE, LCVE, MCCE, MCGE, MCVE, MLCE, MLGE, MLVE, SCCE, SCGE, or SCVE. If a fund lacks a Lipper code, we use Strategic Insight objectives (AGG, GMC, GRI, GRO, ING, or SCG). If neither the Strategic Insight nor Lipper objective is available, we rely on the Wiesenberger Fund Type Code, selecting funds with objectives G, G-I, AGG, GCI, GRI, GRO, LTG, MCG, or SCG. If none of these objectives is available, we retain a fund if it follows a CS policy (i.e., primarily holds common stocks).

[5]We exclude funds with any of the following text strings in their name: "INDEX", "Index", "IDX", "Idx", "S&P", "s&p", "Fixed", "FIXED", "TAX", "tax", "Tax", "CONVERTIBLE", "Convertible", "annuity", "ANNUITY", "ANN", "ann", "VAR", "Var", "CONV", "Conv".

aggregated across mutual funds to compute the stock active shares. Benchmark index returns are also sourced from Refinitiv/Datastream.

Additionally, we acquire data on the Fama-French five factors (market, size, value, investment, and profitability), as well as the momentum factor, from Prof. Kenneth French's website, along with the risk-free interest rate.

Figure 1 provides a comprehensive view of the composition of our dataset over time. Panel I illustrates the evolution of the number of stocks held by mutual funds across the S&P500 (Large-Cap) and S&P400 (Mid-Cap) indices. In the S&P500 (Panel I.A), mutual funds consistently increased their holdings of large-cap stocks, starting with above 350 stocks in the early 2000s, despite some notable fluctuations. By the early 2020s, mutual funds held over 450 stocks from the S&P500, indicating stable and substantial exposure to large-cap equities. In the S&P400 (Panel I.B), holdings followed a similar upward trajectory, though fluctuations in the early 2000s were larger, and recent years saw a slight decline, suggesting a reduced emphasis on mid-cap stocks in mutual fund portfolios. These trends indicate that mutual funds have progressively expanded their exposure across all two indices, with the most consistent growth observed in the S&P500, while mid-cap stocks exhibited greater variability in mutual fund holdings.

Panel II illustrates the percentage of index capitalization for stocks held by mutual funds in the S&P500 and S&P400 indices. In Panel II.A, S&P500 stocks held by mutual funds consistently represent approximately 80% of the index capitalization, with notable stability after 2010. In Panel II.B, the percentage of index capitalization for S&P400 stocks held by mutual funds exhibits more volatility in the beginning of the sample period, but stabilized around 75% in 2010, before increasing steadily to approach 95% by the end of the sample period. These patterns suggest that mutual funds exhibit strong and stable coverage and exposure to large-cap stocks, while showing increasing interest and investment in midcap stocks, particularly after 2010.

Panel III of Figure 1 illustrates the number of mutual funds displaying their holdings for stocks within the S&P500 and S&P400 indices over time.[6] In Panel III.A, the number of funds

[6]We smoothed the data by calculating the three-month rolling average. Otherwise, the data would display high fluctuations and seasonality, with the number of mutual funds sharply increasing at the end of quarters.

disclosing their holdings of S&P500 stocks increases sharply from 2001 to 2005, likely due to the implementation of the SEC's Rule IC-26372 on mutual fund holding disclosure, peaking at over 1,250 funds, and remains relatively stable thereafter. A similar trend is observed in Panel III.B for the S&P400 index, where the number of funds disclosing their holdings also stabilizes above 1,200 funds after a period of rapid growth. A sharp decline towards the end of the sample period (2023) is visible across the S&P500 and S&P400 indices, possibly indicating consolidation within the mutual fund industry or changes in disclosure practices.

2.2 Stock Active Share: Motivation, Definition, and Summary Statistics

2.2.1 Motivating and Defining Stock Active Share

Active mutual funds outperform their benchmark index by selecting winning stocks or timing the market through increased exposure to specific market factors, such as overweighting particular sectors (Cremers and Petajisto; 2009; Petajisto; 2013). These funds tend to have a higher active share, indicating substantial deviations from the benchmark regarding stock weightings. In contrast, closet index funds display a lower active share, closely tracking their benchmark. Since a sufficiently high active share correlates with mutual fund outperformance, the stocks contributing significantly to active share may also drive this outperformance. Therefore, a concentrated portfolio of such stocks will likely outperform the benchmark index.

However, the original active share measure has faced several critiques. Jones and Mo (2020) find that, after the academic publications emphasizing active share's effectiveness as a performance predictor, the measure has become noisier and no longer reliably predicts performance. Our measure captures a different view than the original active share definition, making it distinct and potentially less susceptible to the previously identified market learning effects.[7]

We construct a SAS portfolio comprising stocks that contribute the most to the aggregate active share of mutual funds within a standard benchmark index. Each month, for every stock in the

[7] If this is not the case, the data will reveal it, especially as our sample, beginning in 2007, broadly covers the period following the 2009 publication.

benchmark, we calculate its active share by summing the absolute differences between its weight in each mutual fund portfolio and its weight in the benchmark. Importantly, we measure mutual fund deviations not from specific benchmarks but from a common, standardized benchmark index. The SAS portfolio is then formed based on these stock-level active share measures.

Formally, let $w_{i,t}$ be the weight of stock i in the benchmark index at time t, and $w_{i,j,t}$ be the weight of stock i in mutual fund j's portfolio at time t. Let $N_{i,t}$ represent the number of mutual funds that have disclosed their holdings at time t in stock i. We define the following stock-level measures of active management among mutual funds:

- The **Stock Active Share (SAS)** aggregates the Absolute Deviation From Benchmark (ADFB):

$$\mathrm{SAS}_{i,t} \equiv \sum_{j=1}^{N_{i,t}} \mathrm{ADFB}_{i,j,t} \quad \text{where} \quad \mathrm{ADFB}_{i,j,t} = \left| w_{i,j,t} - w_{i,t} \right| , \tag{1}$$

- The **Net Stock Active Share (NSAS)** aggregates the Deviation From Benchmark (DFB):

$$\mathrm{NSAS}_{i,t} \equiv \sum_{j=1}^{N_{i,t}} \mathrm{DFB}_{i,j,t} \quad \text{where} \quad \mathrm{DFB}_{i,j,t} = \left(w_{i,j,t} - w_{i,t} \right) . \tag{2}$$

We use a summation of ADFB or DFB instead of an average because we aim for our measure to capture the broad market consensus.[8] Using an average (e.g., Jiang et al.; 2014) would, for instance, treat a stock held by a single mutual fund with a specific deviation from the benchmark the same as a stock held similarly by all mutual funds. The latter, however, represents a market consensus, while the former does not. Using the sum, we can clearly differentiate between these two cases, with the stock aligned with the market consensus exhibiting a higher SAS or NSAS value.

2.2.2 Descriptive Statistics on Deviation from Benchmark

Table 1 provides a detailed analysis of mutual fund portfolio weights compared to benchmark indices over the period from August 2007 to December 2023. The table is organized into two

[8] One could consider weighting the deviations by the fund's assets. However, this would give more importance to funds with greater total net assets (TNA), leading to an allocation strategy that aligns more closely with their views.

panels, corresponding to the S&P500 (Panel I) and S&P400 (Panel II). For each benchmark index, sub-panels present summary statistics as follows: sub-panel A displays the weight of a stock disclosed by the mutual fund in its portfolio, while sub-panel B shows the weight of the stock if it were to match its weight in the benchmark index, referred to as the benchmark-matching weight. Sub-panel C reports the stock deviation from the benchmark (DFB), calculated as the difference between the two preceding measures, and sub-panel D provides the absolute deviation from the benchmark (ADFB), which represents the absolute value of the DFB.

The table provides summary statistics for a stock in mutual fund portfolios, aggregated across all funds, all stocks, and all time periods. Each statistic characterizes the typical allocation percentage that mutual funds have invested in a stock, averaged over time and across different funds. This analysis helps to understand the degree of deviation from the benchmark index and the extent of active management by mutual funds.

Across all two indices, mutual funds exhibit a clear tendency to overweight specific stocks compared to the benchmark, particularly in the case of BigCap stocks. On average, mutual funds allocate 0.85% to a typical large-cap stock (with a standard deviation of 1.03%), significantly higher than the benchmark-matching allocation of 0.31% (with a standard deviation of 0.51%). For BigCap stocks, the average ADFB is 0.63% compared to 0.53% for the DFB, with approximately 73.66% of DFBs being positive. In contrast, fewer than half of the deviations for MidCap stocks are positive, indicating that mutual funds tend to underweight these stocks relative to the benchmark. However, the magnitude of positive deviations consistently exceeds that of negative deviations across all indices, reflecting a preference for concentrated positions in certain stocks. The strong correlations (around 0.90) between DFB and ADFB for all indices indicate that absolute deviations are closely aligned with directional deviations, highlighting the substantial active management strategies employed by mutual funds.

2.2.3 Descriptive Statistics on Stock Active Share

Table 2 presents the summary statistics for the same measures introduced in the sub-panels of Table 1, but aggregated across mutual funds. Thus, the reported summary statistics are based on stock-month observations. Similar to Table 1, it is organized into two panels corresponding to the S&P500 (Panel I) and S&P400 (Panel II). For each benchmark index, sub-panel A displays the cumulative weight of a stock disclosed by mutual funds in their portfolios, while sub-panel B shows the cumulative weight of the stock if it were to match its weight in the benchmark index, referred to as the cumulative benchmark-matching weight. Sub-panel C reports the cumulative stock deviation from the benchmark, or Net Stock Active Share (NSAS), calculated as the difference between the two preceding measures, and sub-panel D provides the cumulative absolute deviation from the benchmark, or Stock Active Share (SAS).

Panel I shows that mutual funds allocate 110.34% of their capital to stocks on average, with a standard deviation of 174.33%, compared to a benchmark-matching mean of 40.84% and a standard deviation of 137.27%. This indicates a general tendency for mutual funds to overweight certain stocks relative to the benchmark. Only 0.89% of stock-month observations have negative NSAS values, mainly from large firms like Microsoft, ExxonMobil, and Walmart. The high correlation (0.78) between NSAS and SAS suggests that both measures yield similar results in stock selection, with an even higher correlation (0.94) in the $\mathrm{NSAS} > 0$ subsample.

In Panel II, mutual funds allocate an average of 34.60% to MidCap stocks, with a standard deviation of 31.88%, while the benchmark-matching mean is 20.60%. About 15.89% of NSAS values are negative, indicating more frequent underweighting in MidCap stocks compared to BigCap stocks. The correlation between NSAS and SAS is 0.69, showing a moderately strong relationship.

Overall, mutual funds allocate more capital to BigCap stocks, followed by MidCap, with a higher tendency to underweight MidCap stocks. The high correlations between NSAS and SAS across all indices suggest these measures can be used interchangeably in most stock selection strategies.

3 Preliminary Results on SAS Portfolio Performance and Composition

In this section, we empirically evaluate the performance of an investment strategy that selects, at each period, the stocks with the highest SAS values, repeating this process at the start of each period. The core idea is that one can only outperform the benchmark index with a portfolio that deviates sufficiently from it. Therefore, investing in a limited number of stocks with sufficiently high SAS values, based on mutual fund holdings disclosures, could help achieve this goal. We implement this strategy by sorting the SAS values of the benchmark index constituents, reported as held by at least one mutual fund, into quantiles at each period. We then examine the performance of portfolios composed of assets in each quantile. Intuitively, the lower quantile portfolio, composed of assets with minimal deviations from the benchmark index, should perform similarly to the index. In contrast, the upper quantile portfolios, composed of assets with more significant deviations, are expected to outperform the index. We will test this hypothesis in the following sections.

Let $\mathcal{Q}_{n,t}(\mathrm{SAS})$ denote the nth quantile portfolio at time t, formed by sorting stocks based on their SAS values (assuming these values are known at time t). The return of this portfolio from t to $t+1$ is computed as follows:

$$r_{\mathcal{Q}_{n,t}(\mathrm{SAS}),t+1} = \sum_{i \in \mathcal{Q}_{n,t}(\mathrm{SAS})} \lambda_{i,t} r_{i,t+1} \tag{3}$$

where $r_{i,t+1}$ is the return of stock i from t to $t+1$, and $\lambda_{i,t}$ is the weight of stock i at time t. The weights are normalized such that $\sum_{i \in \mathcal{Q}_{n,t}(\mathrm{SAS})} \lambda_{i,t} = 1$. Stocks in the quantile portfolios can be equally weighted, value-weighted, or weighted according to their SAS values.[9]

Table 3 presents performance metrics for the benchmark indices and SAS-weighted decile, quintile, and tercile portfolios from October 2007 to December 2023, where assets are sorted by SAS values and rebalanced quarterly.[10] These portfolios are called SAS-Oracle,[11] assuming

[9] Subsequent unreported results show that the weighting scheme has minimal impact on the performance of the SAS-sorted portfolio strategies.

[10] We also tested value-weighted and equal-weighted portfolios with similar results, which are omitted for brevity. Monthly rebalancing results are available in the external appendix, showing similar patterns.

[11] The term "Oracle" suggests that investors can perfectly anticipate SAS values, reflecting non-public information

contemporaneous or perfectly anticipated SAS values. The table is divided into two panels: Panel I for the S&P500 and Panel II for the S&P400.

Across all indices, the top quantile portfolios—decile, quintile, or tercile—consistently outperform the lower quantiles and the benchmark index in terms of mean return, standard deviation, Sharpe ratio, and alpha. In Panel I, the top decile portfolio (Q10) delivers a mean return of 1.56%, compared to 0.86% for the benchmark and 0.82% for the bottom decile (Q1). This trend is mirrored in the S&P400, where Q10 outperforms the lower quantiles and the benchmarks. The top portfolios also exhibit lower volatility than the benchmark, with Q10 in the S&P500 having a standard deviation of 4.24%, compared to 4.66% for the benchmark and 7.10% for Q1.

Sharpe ratios rise from the bottom to the top quantiles, with Q10 in the S&P500 achieving 0.35, significantly higher than the benchmark's 0.17 and Q1's 0.10. Similar upward trends are observed across quintile and tercile portfolios, further confirming the superior risk-adjusted performance of the top quantile portfolios. Kurtosis improves across the quantiles, and though skewness becomes more negative in top portfolios, the adjusted Sharpe ratio remains higher for top quantiles due to better overall performance.

Alpha values, which measure risk-adjusted returns against the Fama-French five factors and the momentum factor, further reinforce the outperformance of top quantile portfolios. In the S&P500, Q10 generates an alpha of 0.74%, compared to -0.02% for the benchmark and -0.13% for Q1. Similar patterns are observed in the S&P400, where Q10 generates an alpha of 1.41%, compared to -0.07% for the benchmark and 0.27% for Q1.

Although quintile and tercile portfolios show similar performance patterns, their Sharpe ratios and alphas are generally lower than those of the top decile portfolios, suggesting that focusing on a smaller set of high-SAS stocks enhances portfolio efficiency.

Overall, the SAS strategy consistently identifies high-performing stocks, with top portfolios outperforming both the benchmark indices and lower quantiles regarding returns and risk-adjusted performance. Concentrating on the top decile portfolios allows for significant outperformance

held by mutual fund managers.

without increasing risk, underscoring the value of active mutual fund managers' stock selection.

A long-term evaluation of the SAS-Oracle portfolios from October 2007 to December 2023 shows significant outperformance over passive strategies. The top decile SAS-Oracle portfolio achieved cumulative returns of approximately 300% for BigCap stocks (and over 400% for MidCap stocks), translating to a fourfold increase in initial investment. For example, a $1,000 investment in the BigCap SAS-Oracle portfolio would have grown to $3,850 by December 2023, compared to a 147% return for the S&P500. The strategy also delivers superior risk-adjusted returns (alphas), with cumulative alphas of around 130% for BigCap stocks and 250% for MidCap. In contrast, benchmark indices' alphas remain close to zero, indicating passive returns are largely driven by systematic risk exposure.[12]

The SAS-Oracle strategy finally shows strong persistence and consistency in selecting high-performing stocks, particularly for the BigCap. In this category, Microsoft is selected 98.5% of the time in the SAS-Oracle top ventile portfolio, followed by JP Morgan Chase (89.7%), Alphabet (78.5%), Visa (61.5%), and Meta (57.4%). For MidCap stocks, Reinsurance Group of America leads with a 44.6% selection rate, followed by ANSYS (30.3%), IDEX (16.9%), Carlisle (15.9%), and BJS Wholesale Club (14.4%).[13]

This section provides empirical evidence that mutual fund holdings contain collective information capable of consistently outperforming the market when effectively utilized. This suggests that the success of the SAS strategy is driven by managerial skill, not random chance. If fund managers lacked skill, the SAS strategy's performance would be random and fail to surpass the benchmark. Instead, the strategy aggregates their expertise into superior knowledge, similar to how artificial intelligence combines information efficiently. The persistence of certain stocks in the SAS-Oracle portfolio further indicates that the strategy is not random but potentially predictable. Although delays in fund holdings disclosures limit real-time application, this persistence offers valuable insights for stock selection predictability, which will be explored in the next section.

[12] Details on cumulative returns and alphas of the SAS-Oracle strategy and benchmark indices are provided in Figure B1 in the external appendix.

[13] For an illustration of stock selection history, refer to Figure B2 in the external appendix.

4 Predicting Managers' Convictions About Stock Values

Following DeMiguel et al. (2023), we apply various machine learning methods to predict stocks' SAS measures and select stocks for our investment portfolio. This involves using the history of mutual fund holdings and stock characteristics to learn decision rules that mimic the skills of mutual fund managers. The objective is to build a model that predicts future SAS values based on current stock characteristics, akin to how managers select stocks using publicly available information.

We frame this supervised learning task as either a regression problem for predicting the stock's SAS measure or a classification problem for predicting stock selection into the SAS strategy. Simple models, such as Ordinary Least Squares, serve as baselines, while more advanced models include Elastic Net, Random Forests, Gradient Boosting, and Deep Neural Networks.[14]

Ordinary Least Squares (OLS) assume a linear relationship between predictors and outcomes but may lead to overfitting, especially with many predictors. Elastic Net (ENET) addresses this by adding regularization to prioritize important predictors while penalizing less significant ones. Random Forests and Gradient Boosting capture non-linear relationships and interactions between predictors. While Random Forests (RF) reduce variance by averaging predictions across multiple decision trees, Gradient Boosting (GB) iteratively improves predictions by focusing on observations poorly predicted by previous trees. Lastly, we employ Deep Neural Networks (DNN) to approximate complex functions through layers of interconnected neurons.

4.1 Data and Procedure

We utilize 90 stock-level variables, as detailed in Table 4, following Gu et al. (2020), alongside 9 macroeconomic predictors described in Welch and Goyal (2008). These include key variables such as the dividend-price ratio, earnings-price ratio, book-to-market ratio, net equity expansion,

[14] These models are trained and validated using cross-validation to ensure they generalize well to unseen data, and we use a modified version of the code provided by DeMiguel et al. (2023) to verify our calculations, except for Deep Neural Networks, which are not considered by these authors. Further details about the mathematical formulations of these models, including regularization techniques, decision tree splitting, and neural network architecture, can be found in the external appendix.

Treasury-bill rate, term spread, default spread, and stock variance.[15]

Data pre-processing involves replacing missing values with cross-sectional averages for each period. Additionally, all explanatory variables are standardized by subtracting their cross-sectional mean and dividing by the cross-sectional standard deviation on a monthly basis.

The dataset comprises stock-month observations, split into training, validation, and test samples. A recursive training procedure is applied with an expanding estimation window. The model is updated quarterly, incorporating newly available data. Our initial training period spans from August 2007 to December 2017, with the final 12 months of this period reserved for validation.[16] The model parameters are updated at quarterly as the portfolio's rebalancing frequency.

We optimize model hyperparameters—such as regularization parameters in Elastic Net, the number of trees in decision trees, and architecture specifics for deep neural networks (e.g., layers, neurons, activation functions, regularization rates, and learning rates)—using the validation set. The aim is to minimize overfitting and improve the model's ability to generalize to unseen data. We base our hyperparameter selection on minimizing the loss function and ensuring convergence in performance between the training and validation datasets. Successful convergence indicates the model's ability to predict future values in the test set.

We evaluate the predictive performance of the models both statistically and financially. Statistically, we use two metrics: the traditional out-of-sample R^2 ($R^2_{OOS_T}$) and a modified out-of-sample R^2 ($R^2_{OOS_M}$), designed to account for potential biases in individual stock SAS predictions. These metrics are defined as:

$$R^2_{OOS_T} = 1 - \frac{\sum_{(i,t)\in\mathcal{T}}(SAS_{i,t} - \widehat{SAS}_{i,t})^2}{\sum_{(i,t)\in\mathcal{T}}(SAS_{i,t} - \overline{SAS}_{\mathcal{T}})^2}, \tag{4}$$

$$R^2_{OOS_M} = 1 - \frac{\sum_{(i,t)\in\mathcal{T}}\left((SAS_{i,t} - \overline{SAS}_{\mathcal{T}}) - (\widehat{SAS}_{i,t} - \overline{\widehat{SAS}}_{\mathcal{T}})\right)^2}{\sum_{(i,t)\in\mathcal{T}}(SAS_{i,t} - \overline{SAS}_{\mathcal{T}})^2} \tag{5}$$

[15] We obtain monthly updates for these variables from Amit Goyal's website.

[16] Details of the recursive training procedure and additional methodological clarifications are available in the external appendix.

Here, $\mathcal{T}$ represents the stock-time observations in the test sample.

4.2 Prediction Results

This section presents the results of the machine-learning (ML) models used to predict SAS measures. We recall the motivations behind our approach. First, there is a delay in the release or unavailability of real-time information about mutual fund portfolio holdings. Second, we aim to reverse engineer mutual fund managers' stock-picking decisions to learn from their stock selection skills.

Table 5 presents the out-of-sample prediction performance measures across the various machine learning models we consider. The dependent variable is the logarithm of the stock active share, with the predictors being macro variables and stock characteristics. The table is organized into two panels: Panel I presents the results for the S&P500 index and Panel II covers the S&P400 index. Each panel is organized into three sub-panels: decile portfolios (.A), quintile portfolios (.B), and tercile portfolios (.C), which correspond to different quantile groupings of stocks based on the predicted variable. The *DepVar* (dependent variable) shows the actual values, while *PredVar* (predicted variable) reflects the values predicted by each machine learning method. Other columns include the coefficient of variation (*CV*), and two versions of out-of-sample R-squared metrics: the traditional out-of-sample $R^2_{OOS_T}$ and the modified out-of-sample $R^2_{OOS_M}$, which accounts for bias in predicting individual stock SAS.

Across the benchmark indices, the coefficient of variation (*CV*) tends to decrease as we move up the quantiles. It is also lower for BigCap stocks reflecting for stock active shares what is commonly recognized for stock return volatility. This stylized fact can be explained by several considerations. On the one hand, BigCap stocks are generally more liquid and widely held, which stabilizes demand and reduces the variability of SAS. Additionally, while SAS measures the demand for stocks by institutional investors, it is indirectly linked to stock prices, as increased or concentrated demand for a specific stock can potentially drive its price upward, and vice versa.

Focusing on Panel I, the results show consistency across different types of quantiles and machine-learning methods. The modified R^2 is generally positive, although it is negative for some quantiles and models. For example, in decile portfolios, R^2 is negative for Q1 with GB, OLS, and Elastic Net and for Q10 with OLS and Elastic Net. Despite these exceptions, models like DNN, RF, and GB generally perform better. Regardless of the machine learning method, the R^2 values tend to increase gradually from the lower to the higher quantiles. For instance, in decile portfolios, the modified R^2 with RF ranges from 0.13 (Q1) to 0.34 (Q10), while with GB it ranges from 0.13 (Q2) to 0.30 (Q10). OLS and Elastic Net show weaker performance, with R^2 values ranging from 0.05 (Q2) to 0.20 (Q9) for OLS and 0.06 (Q2) to 0.21 (Q9) for Elastic Net. DNN shows the strongest performance, with R^2 values increasing from 0.03 (Q1) to 0.36 (Q10).

The same pattern is observed in quintile portfolios, where the top quintile (Q5) displays modified R^2 values of 0.36 (RF), 0.34 (GB), and 0.37 (DNN), and tercile portfolios, where the top tercile (Q3) shows values of 0.39 (RF), 0.36 (GB), 0.16 (OLS and Elastic Net), and 0.38 (DNN). The traditional R^2 values are generally lower than the modified R^2, as the modified version corrects for bias in the prediction. Still, they follow a similar pattern across quantiles and ML methods.

Regarding proximity between actual SAS and predicted SAS, OLS and Elastic Net tend to show the closest values in the top quantile portfolios. For example, the difference between the actual and predicted values in Q10 is 0.06 for OLS and Elastic Net, compared to 0.43 for Random Forest, 0.24 for Gradient Boosting, and 0.31 for DNN.

Another important observation is the improvement in R^2 values and the reduction in prediction bias as the portfolio concentration decreases, i.e., when more assets are included in the portfolio. For instance, with Random Forests, the modified R^2 increases from 0.34 in the top decile to 0.36 in the top quintile and 0.39 in the top tercile. When all assets are included in a global portfolio, the R^2 is 0.48. Similarly, the prediction bias, which can be approximated by the difference between the modified and traditional R^2, decreases from 0.27 in the top decile to 0.15 in the top quintile and 0.09 in the top tercile and is just 0.03 for the global portfolio. These improvements are even more pronounced with DNN, where the bias decreases from 0.13 in the top decile to 0.01 in the top tercile.[17] In the global portfolio, the bias disappears entirely, as both R^2 values equal 0.43.

[17] While the top decile portfolio was identified as the best candidate for implementing the SAS-Oracle strategy, offering greater concentration with similar performance metrics to the top quintile and tercile portfolios, the recom-

Panel II of Table 5 shows that the predictive quality for MidCap stocks is generally lower than for S&P500, with small or negative modified R^2 values across most quantiles and methods. However, DNN and RF outperform, achieving modified R^2 values of 13% and 7% for the top decile, 11% and 9% for the top quintile, and 10% for the top tercile in both cases.

Gu et al. (2020) and Kaniel et al. (2023) also provide out-of-sample R^2 to assess the performance of their machine learning specifications in predicting the monthly risk-adjusted returns out-of-sample, for stocks and mutual funds, respectively. The R^2 values achieved for the whole sample range from -3.46% to 0.40% for the former and from -1.60% to 5.00% for the latter.[18] While these authors predict returns, we predict stock active shares, which reflect managers' convictions about asset values. Our predictive performance is superior in this context. In comparison, this suggests that stock active shares may provide more reliable insights into managerial expertise and stock selection than direct return predictions.

The predictions of SAS by machine learning methods are highly correlated, as shown in Figure 3, with correlations ranging between 0.75 and 0.89 when applied to S&P500 stocks. These correlations vary between 0.71 and 0.82 for MidCap stocks.

Since the quantile portfolios in Table 5 are based on predicted SAS values, the *PredVar* column should increase from the bottom to the top quantile, which is consistently observed across methods. Similarly, the actual SAS values in the *DepVar* column show the same upward trend, confirming that the SAS-feasible strategy aligns with the SAS-Oracle strategy. High correlations between actual and predicted SAS values—ranging from 0.59 to 0.66 for BigCap, and 0.41 to 0.49 for MidCap—further support the reliability of these predictions (Figure 3). These *out-of-sample* results demonstrate model generalizability, reinforcing the SAS-feasible strategy, whose performance is compared to the SAS-Oracle and other strategies in the next section.

mendation is more nuanced for the SAS-feasible strategy. Less concentrated portfolios, such as quintile or tercile, may be preferred due to improved SAS predictability and reduced bias.

[18] See Table 1 in Gu et al. (2020) and Table 3 in Kaniel et al. (2023).

5 Performance of the Machine Learning-based SAS Strategy

The SAS-feasible strategy uses machine learning to form portfolios by ranking stocks based on their predicted SAS measure. The approach involves going long on the top-quantile portfolio of stocks with the highest predicted SAS values. We use predictors measured at time t and a machine learning model trained on data up to period t to anticipate the time-t stock active share. The strategy's financial performance is evaluated out-of-sample using real-time data, with performance metrics including mean, standard deviation, skewness, kurtosis, Sharpe ratio, and adjusted Sharpe ratio. We also report the average risk-adjusted return (alpha) relative to the Fama-French and momentum factors, with standard errors calculated using the Newey-West HAC variance-covariance matrix to account for heteroskedasticity and autocorrelation.

5.1 The SAS-feasible Investment Strategy with Quarterly Rebalancing

Table 6 provides summary statistics for the SAS-feasible investment strategy with quarterly rebalancing. The table is organized into three blocks representing top-quantile portfolios (Top 10, Top 5, Top 3) and five machine-learning methods (RF, GB, OLS, ENET, DNN), across two panels corresponding to the benchmark indices (S&P500 and S&P400). For comparison, the table also includes performance statistics for analyst recommendations (AR), the SAS-Oracle strategy (ORA) based on actual SAS measures, and the benchmark index. The out-of-sample evaluation covers the period from January 2018 to December 2023.

For all indices, the SAS-feasible strategy (Top 10, Top 5, Top 3 portfolios) consistently delivers lower average returns than the SAS-Oracle portfolios but still significantly outperforms the benchmark index. For example, in the S&P500, the Top 10 DNN portfolio achieves an average return of 1.39%, compared to 1.42% for the SAS-Oracle portfolio and 1.09% for the benchmark. This trend is consistent across the S&P400 index, with SAS-feasible portfolios trailing the SAS-Oracle in returns but comfortably exceeding the benchmark.

Regarding volatility, the SAS-feasible portfolios generally show lower or comparable volatility to the benchmark index and the SAS-Oracle portfolios. For instance, in the S&P500, the Top

10 DNN portfolio has a volatility of 5.02%, lower than both the SAS-Oracle (5.70%) and the benchmark (5.21%). Similarly, across the S&P400 index, the SAS-feasible portfolios exhibit moderate volatility, demonstrating better risk control compared to the benchmark.

Sharpe ratios indicate that the SAS-feasible strategy consistently delivers better risk-adjusted performance than the benchmark across all indices, closely approaching that of the SAS-Oracle strategy. In the S&P500, the Top 10 DNN portfolio achieves a Sharpe ratio of 0.25, outperforming both the benchmark's 0.18 and the SAS-Oracle's 0.22. A similar trend is observed in the S&P400 index, where the SAS-feasible strategy outperforms the benchmark and remains competitive with the SAS-Oracle strategy. This demonstrates that the SAS-feasible strategy provides strong risk-adjusted returns, even with predicted rather than actual SAS values.

Alpha values reinforce the profitability of the SAS-feasible strategy, with consistently positive alphas across all indices and portfolios. In the S&P400, for instance, the Top 10 DNN portfolio achieves an alpha of 0.68%, surpassing the AR's 0.18% and the benchmark's -0.06%. This pattern is similarly observed in the S&P500, where SAS-feasible portfolios generate significant alphas, rivaling the SAS-Oracle portfolios and outperforming the benchmark indices.

We compare the performance of our SAS-feasible strategy with other machine learning-based strategies for investing in stock or mutual fund portfolios. Gu et al. (2020) report annualized Sharpe ratios ranging from -0.23 in the lowest decile to 0.81 in the highest decile for prediction-sorted stock portfolios. Kaniel et al. (2023) report monthly Sharpe ratios from -0.23 to 0.15 for mutual fund portfolios, and DeMiguel et al. (2023) show a monthly Sharpe ratio of up to 0.192 in a long-only mutual fund portfolio. These comparisons demonstrate that the performance of our SAS-feasible strategy is competitive with existing machine learning-based investment methods.

In summary, while the SAS-Oracle strategy yields slightly higher returns, the SAS-feasible strategy still delivers superior returns and risk-adjusted performance compared to the benchmark. The narrow performance gap between the SAS-feasible and SAS-Oracle portfolios highlights the ML models' robust learning ability, allowing them to closely replicate the performance of the otherwise impractical SAS-Oracle strategy. SAS-feasible portfolios often have lower standard deviations, indicating better risk management, and consistently outperform the benchmark regarding Sharpe ratios and alphas. The significant positive alphas suggest the strategy remains profitable, even after considering transaction costs, which are likely low due to quarterly rebalancing. Figure 2 shows the sustained performance of machine learning-based portfolios over time.

Table 7 provides additional performance measures, further confirming the analysis presented earlier. For instance, the comparison of the Top 10 DNN, Top 10 ORA, and Index portfolios across the S&P500 andS&P400 indices shows that the DNN model consistently outperforms the benchmark index in all panels, demonstrating superior risk-adjusted returns (Sortino, Information Ratio) and better risk management (lower Drawdown and VaR). In the S&P500, the DNN portfolio surpasses the ORA strategy, delivering higher Sortino and Information Ratios while maintaining lower risk exposure. However, in the S&P400, while the ORA portfolio achieves better Sortino and Information Ratios, the DNN portfolio still provides more robust risk management with a lower Drawdown. These results highlight the machine-learning models' ability to deliver competitive returns and manage risk effectively, rivaling the infeasible SAS-Oracle strategy while consistently outperforming the passive strategy across all benchmark indices.

5.2 SAS-feasible versus Analyst Recommendation Consensus

Analysts serve as information intermediaries, gathering, analyzing, and producing investment-related insights for the broader community (Kothari et al.; 2016). Unlike mutual fund managers, who trade on behalf of their clients and disclose their holdings ex-post to comply with regulations, analysts provide ex-ante recommendations, allowing investors to act immediately on their advice. Both analysts and fund managers bear reputational and financial risks, motivating them to exert considerable effort in identifying valuable stocks. Thus, comparing the performance of portfolios formed based on machine learning predictions of mutual fund holdings (SAS-feasible strategy) and those based on analyst recommendations (AR strategy) is a worthwhile exercise.

In this section, we evaluate the performance of investment strategies that involve buying top sorted stocks based on consensus analyst recommendations or machine learning-predicted SAS measures derived from mutual fund managers. Analysts regularly provide stock recommendations ranging from 1 (strong buy) to 5 (sell). To facilitate comparison with SAS values, we transform the original recommendation scores by subtracting the value from 6, ensuring that higher values represent stronger stock valuations. This transformation allows a consistent interpretation of stocks sorted based on transformed analyst recommendations or predicted SAS measures.

Figure 3 highlights the correlations between analyst recommendations and ML-predicted SAS measures for different benchmark indices from January 2018 to December 2023. The results indicate low but positive correlations, typically ranging from 0.1 to 0.2. Linear models such as OLS and Elastic Net show slightly higher correlations with analyst recommendations than non-linear machine learning models, with the strongest correlations observed in BigCap stocks. Therefore, our predicted measure of mutual fund managers' attention to stocks aligns with analysts' consensus about stock valuation, though the two measures are far from being perfectly correlated.

We now turn to the investment performance comparison. The AR columns in Table 6 and Table 7 report the performance of the AR strategy.[19] Similar to the SAS-feasible strategy, the AR strategy involves sorting stocks based on analyst consensus recommendations and investing in the top quantile of stocks with the strongest buy recommendations, with quarterly rebalancing. Despite performing well relative to the benchmark indices, the AR strategy falls short compared to the SAS-feasible strategy. The average returns and alphas, as measured with the Fama-French 5 factors and the momentum factor, are significantly higher in the AR portfolio compared to the benchmark index for S&P400. However, for the S&P500, the differences between the AR portfolio and the benchmark index, even though positive, are not statistically significant at the standard level. Across various metrics—average return, standard deviation, and Sharpe ratio— the SAS-feasible and SAS-oracle portfolios consistently outperform the AR portfolios in the S&P400 (see Table 6 and A4). Results are more nuanced for the S&P500 index.

Across benchmark indices, the portfolios formed on lagged SAS (ORA_1, ORA_2, and ORA_3) offer similar performance as AR portfolios. Portfolios constructed using lagged SAS for the

[19]Table A4 in appendix reports the results of tests of differences in mean return, Sharpe ratio, and alphas between the SAS and AR portfolios.

S&P400 index typically underperform relative to oracle SAS and machine learning-based SAS portfolios, while exhibiting performance comparable to AR portfolios. When it comes to the S&P500 index, only Oracle's portfolios do much better than the other strategies, which yield about the same performance results. The comparable performance between the $\mathbf{MLSAS}$ strategy and the AR strategy, together with their similar challenges in surpassing the benchmark index, may be explained by the greater informational efficiency of the BigCap market. Big companies are closely monitored by the public, and any information pertaining to these stocks is rapidly reflected in their prices; this may elucidate the challenge of significantly surpassing the benchmark for these stocks.

Additionally, Table 7 shows that the $\mathbf{MLSAS}$ portfolios, in particular the Top 10 DNN portfolios, surpass the AR portfolios in both performance and risk management across the S&P500 and S&P400 indices. DNN portfolios achieve higher risk-adjusted returns and better control of downside risks, despite slightly higher turnover in some cases. These results demonstrate that machine learning strategies like DNN offer superior portfolio performance and risk management compared to traditional analyst-based approaches.

In summary, while analyst recommendations offer valuable insights, machine learning-based predictions of SAS measures derived from mutual fund managers provide superior investment guidance, consistently yielding higher returns. However, this analysis highlights the ongoing importance of analysts in delivering meaningful insights. Given the expertise required to surpass analyst-driven strategies, it may still be worthwhile for investors to pay for access to high-quality analyst advice. Our findings also underscore the advantages of active management, particularly in the selection of MidCap stocks, which may exhibit lower informational efficiency relative to BigCap stocks.

5.3 Which covariates matter?

We now examine the inner workings of the machine learning models to investigate the relative importance of the input variables driving model performance. We use Local Interpretable Model-agnostic Explanations (LIME), as introduced by Ribeiro et al. (2016). The objective of LIME is to identify an interpretable model that is locally faithful to the machine learning model's prediction function and understandable to humans, regardless of the features used by the model. LIME identifies which features are most important in explaining individual predictions by approximating the model locally with a simpler, interpretable representation. Additionally, it shows the direction of each feature's influence, revealing whether it contributed favorably or unfavorably to the predicted value for a specific entity.

By aggregating explanations across all observations in the test sample, we can determine the variables that primarily drive the model predictions. LIME's detailed feature importance information enables better comprehension of the machine learning models' predictions, builds user trust, and can be leveraged to improve model performance or communicate more effectively with clients about portfolio outcomes.

Figures 4 illustrate the most important features the machine learning models identified following quarterly retraining. The predictors are ranked such that the top 10 most important features are listed from most to least important. Red bars indicate that a characteristic positively contributes to the predicted value, while dark blue bars indicate negative contributions. The models show a high level of agreement on the critical factors influencing stock selection, reinforcing the reliability of the predictions. The key features can be broadly categorized into three groups. The first group relates to past/recent performance measures, including the stock's last period SAS value, momentum over 1, 12, or 36 months, and past Treynor and Sharpe ratios. These variables generally capture momentum effects, as past top performers are expected to continue performing well. However, reversal effects may appear when past top performers are predicted to underperform in the future.

The second group comprises macroeconomic predictors, such as the treasury bill rate (*tbl*), default yield spread (*dfy*), book-to-market value of DJIA (*bm*), net equity expansion (*ntis*), and term spread (*tms*). Both *tbl* and *tms* negatively predict future SAS values, suggesting that stocks favored by managers tend to underperform when interest rates rise or term spreads expand. In contrast, variables like *bm*, earnings-to-price ratio (*E12*), and *dfy* positively contribute to the predicted SAS measure, indicating that stocks appreciated by fund managers are likely to benefit when these factors increase.

The third group includes signals of stock growth opportunities and risks, represented by variables such as the convertible debt indicator (*convind*), market capitalization (*mc*), long-term debt growth (*lgr*), Sin stock indicator (*sin*), 36-month rolling Sharpe Ratio (*sr36m*), industry-adjusted change in asset turnover (*chatoia*), accruals volatility (*stdacc*), and sales-to-cash ratio (*salecash*). Consistent with Gu et al. (2020), sin stocks appear unattractive to fund managers, as indicated by their negative contribution to the predicted SAS measure.

6 Carbon Emissions and SAS Investment

Responsible investing has increasingly gained importance in investor performance criteria. However, the question of whether responsible investing, particularly with respect to carbon emissions, impacts financial performance remains open.[20] To address this issue, we extract carbon emissions data for stocks in the benchmark indices (S&P400 and S&P500) from the Eurofidai database using firm ISIN codes. We compute carbon emissions intensity by dividing the sum of scope 1 (direct) and scope 2 (indirect) CO_2 emissions by the firm's market capitalization. A higher emission intensity indicates that a firm produces more carbon emissions per dollar of market capitalization, either directly or indirectly.[21] Bolton and Kacperczyk (2023) and Aswani et al. (2023) provide convincing arguments for the use of emission level and emission intensity respectively in studying the market perception of the impact of firm emissions on climate change.[22] Therefore, we also run the same regressions but using carbon emissions level instead of carbon emissions intensity as our main explanatory variable.

[20] Matsumura et al. (2014) find that increase in carbon emissions negatively affects the firm's value for S&P500 firms; meaning that market penalizes firms for their carbon emissions. In contrast, Bolton and Kacperczyk (2021, 2023) find higher stock returns associated with higher levels and growth rates of carbon emissions for the US firms and internationally, which they interpreted as investors demanding compensation for their exposure to carbon emission risk. Furthermore, Aswani et al. (2023) using carbon intensity instead of carbon emissions, find no association between emissions and returns.

[21] The data is winsorized at the 1% level to minimize the effect of outliers.

[22] Bolton and Kacperczyk (2023) argue that : "*What the world needs and aims for is first a reduction in carbon emission levels, and second only an improvement in carbon efficiency.*" Whereas Aswani et al. (2023)'s counterargument is that: "*Emissions arise from a firm's core operations and, absent significant short-term innovations in a firm's production process, unscaled emissions are largely determined by the quantity of goods produced and sold.*

We then examine the relationship between stock carbon emissions and fund manager conviction about stock value, as measured by our SAS metric. We specify the following econometric model:

$$\ln\left(SAS_{i,t}\right) = \beta_0 + \beta_1 \ln\left(Carbon_{i,t}\right) + \gamma_i + \lambda_t + \varepsilon_{i,t} \tag{6}$$

where $SAS_{i,t}$ is the stock i's average SAS measure during year t, and $Carbon_{i,t}$ is either the stock's carbon emissions intensity or carbon emissions level. The terms γ_i and λ_t represent stock and time fixed effects, respectively. β_1 captures the elasticity of a stock's SAS with respect to its carbon emissions.

Panel A (respectively, Panel B) of Table 8 shows the regression results with carbon emissions intensity (respectively, carbon emissions level) as main explanatory variable. The elasticity of the SAS measure to emission intensity is negative and statistically significant, whereas it is non-significant or significantly positive when emission level is used. The negative coefficient of emission intensity suggests that stocks heavily used by fund managers to deviate from benchmark indices and contribute to overperformance also tend to have lower emission intensity. This means that a stock improvement in emission efficiency (a decrease of emission intensity) increases its likelihood to be selected into SAS portfolio. For MidCap and BigCap stocks, the effect is even stronger: a 10% increase in emission intensity results in a 2.9% and 1.5% decrease in their SAS measures, respectively. These shifts in SAS values due to changes in emission intensity are substantial enough to influence the selection or exclusion of stocks from SAS portfolios.

In contrast, for BigCap stocks, a 10% increase in carbon emissions level within a firm leads respectively to 0.6% increase in its SAS measure. The positive coefficient of emission level suggests that higher carbon emissions level predicts higher stocks SAS value and therefore higher stock performance. This interpretation aligns with the view that stocks with higher level of carbon emissions earn a positive carbon premium that reflects exposure to climate related concerns and their implied risks (physical, transition, or regulatory risks) to investors (Bolton and Kacperczyk; 2021, 2023). For MidCap stocks, the effect is not significant.

While we attempt to control for endogeneity using stock and time fixed effects, there remains a potential for bias in the estimated elasticities. Nevertheless, our findings strongly suggest that the SAS strategy is compatible with ESG (Environmental, Social, and Governance) criteria as it reduces carbon intensity of the portfolio, and it can be used to achieve both financial performance and sustainable investment objectives.

7 Conclusion

This study demonstrates that mutual fund managers collectively possess stock-picking skills that outperform passive benchmarks and consensus-based strategies from analyst recommendations. By applying machine learning techniques to mutual fund holdings, investors can effectively leverage these skills. Our analysis shows that constructing a portfolio of stocks with the most significant deviations from the benchmark leads to superior performance, as measured by the adjusted Sharpe ratio, while delivering positive risk-adjusted returns (i.e., alphas) after accounting for systematic risk exposures such as Fama-French five factors and momentum factor. Thus, by analyzing mutual fund stock holdings, we demonstrate that investors can harness fund manager expertise through machine learning.

The evaluation of the stock active share's feasible investment strategy does not account for transaction costs, such as bid-ask spreads, brokerage commissions, or the market impact of trading. While these costs are often overlooked in studies due to their quantification challenges, they play a critical role in the real-world profitability of investment strategies. Future research should prioritize incorporating these factors into performance assessments, as advocated by Ferson (2010).

We also examined the relationship between stock active share and firm carbon emissions level and intensity, finding both a negative correlation for emission intensity and a positive correlation for emission level. This suggests that the financial performance of the stock active share's feasible investment strategy may be compatible with environmental performance. However, we approach this interpretation cautiously, as the observed association may hide various causal effects that warrant further investigations.

References

Agarwal, V., Wei, J., Tang, Y. and Yang, B. (2013). Uncovering hedge fund skill from the portfolio holdings they hide, *The Journal of Finance* **68**(2): 739–783.

Antón, M., Cohen, R. B. and Polk, C. (2021). Best ideas, *Working paper* .

Antón, M., Cohen, R. B. and Silli, B. (2008). Best ideas, *Working paper* .

Aswani, J., Raghunandan, A. and Rajgopal, S. (2023). Are carbon emissions associated with stock returns?*, *Review of Finance* **28**(1): 75–106.

Barber, B., Lehavy, R., McNichols, M. and Trueman, B. (2001). Can investors profit from the prophets? security analyst recommendations and stock returns, *The Journal of Finance* **56**(2): 531–563.

Berk, J. B. and van Binsbergen, J. H. (2015). Measuring skill in the mutual fund industry, *Journal of Financial Economics* **118**(1): 1–20.

Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk?, *Journal of Financial Economics* **142**(2): 517–549.

Bolton, P. and Kacperczyk, M. (2023). Global pricing of carbon-transition risk, *The Journal of Finance* **78**(6): 3677–3754.

Bryzgalova, S., Pavlova, A. and Sikorskaya, T. (2023). Retail trading in options and the rise of the big three wholesalers, *The Journal of Finance* **78**(6): 3465–3514.

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

Cremers, K. J. M. and Petajisto, A. (2009). How active is your fund manager? a new measure that predicts performance, *The Review of Financial Studies* **22**(9): 3329–3365.

Daniel, K., Grinblatt, M., Titman, S. and Wermers, R. (1997). Measuring mutual fund performance with characteristic-based benchmarks, *The Journal of Finance* **52**(3): 1035–1058.

DeMiguel, V., Gil-Bazo, J., Nogales, F. J. and Santos, A. A. (2023). Machine learning and fund characteristics help to select mutual funds with positive alpha, *Journal of Financial Economics* **150**(3): 103737.

Eaton, G. W., Green, T. C., Roseman, B. S. and Wu, Y. (2022). Retail trader sophistication and stock market quality: Evidence from brokerage outages, *Journal of Financial Economics* **146**(2): 502–528.

Fama, E. F. and French, K. R. (2010). Luck versus skill in the cross-section of mutual fund returns, *The Journal of Finance* **65**(5): 1915–1947.

Ferson, W. E. (2010). Investment performance evaluation, *Annual Review of Financial Economics* **2**(1): 207–234.

Green, J., Hand, J. R. M. and Zhang, X. F. (2017). The Characteristics that Provide Independent Information about Average U.S. Monthly Stock Returns, *The Review of Financial Studies* **30**(12): 4389–4436.

Gu, S., Kelly, B. and Xiu, D. (2020). Empirical Asset Pricing via Machine Learning, *The Review of Financial Studies* **33**(5): 2223–2273.

Jiang, H., Verbeek, M. and Wang, Y. (2014). Information content when mutual funds deviate from benchmarks, *Management Science* **60**(8): 2038–2053.

Jones, C. S. and Mo, H. (2020). Out-of-Sample Performance of Mutual Fund Predictors, *The Review of Financial Studies* **34**(1): 149–193.

Kacperczyk, M., Sialm, C. and Zheng, L. (2005). On the industry concentration of actively managed equity mutual funds, *The Journal of Finance* **60**(4): 1983–2011.

Kacperczyk, M., Sialm, C. and Zheng, L. (2006). Unobserved Actions of Mutual Funds, *The Review of Financial Studies* **21**(6): 2379–2416.

Kaniel, R., Lin, Z., Pelger, M. and Van Nieuwerburgh, S. (2023). Machine-learning the skill of mutual fund managers, *Journal of Financial Economics* **150**(1): 94–138.

Kelly, B. and Xiu, D. (2023). Financial machine learning, *NBER Working Paper* (31502): 1035–1058.

Kothari, S., So, E. and Verdi, R. (2016). Analysts' forecasts and asset pricing: A survey, *Annual Review of Financial Economics* **8**(Volume 8, 2016): 197–219.

Li, B. and Rossi, A. G. (2021). Selecting mutual funds from the stocks they hold: A machine learning approach, *Working Paper* .

Matsumura, E. M., Prakash, R. and Vera-Muñoz, S. C. (2014). Firm-value effects of carbon emissions and carbon disclosures, *The Accounting Review* **89**(2): 695–724.

Petajisto, A. (2013). Active share and mutual fund performance, *Financial Analysts Journal* **69**(4): 73–93.

Ribeiro, M. T., Singh, S. and Guestrin, C. (2016). "why should i trust you?": Explaining the predictions of any classifier.

Welch, I. and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction, *The Review of Financial Studies* **21**(4): 1455–1508.

Wermers, R., Yao, T. and Zhao, J. (2012). Forecasting stock returns through an efficient aggregation of mutual fund holdings, *The Review of Financial Studies* **25**(12): 3490–3529.

Yan, X. S. and Zhang, Z. (2007). Institutional investors and equity returns: Are short-term institutions better informed?, *The Review of Financial Studies* **22**(2): 893–924.

Table 1: **Descriptive statistics of fund portfolio weights and deviations from benchmark**
For each benchmark index (S&P500 in Panel I and S&P400 in Panel II), the table displays descriptive statistics of the stock weight disclosed by the mutual fund, the stock weight that would be disclosed if the mutual fund invested the same weight as in the benchmark index, and the deviations from benchmark (DFB and ADFB), considering all mutual funds, all benchmark index constituents with a mutual fund holding disclosure, and across all time periods. Reported statistics are the mean, standard deviation, minimum, percentiles, maximum, and correlation between DFB and ADFB. Each statistic is a single value that describes the typical weight or deviation from benchmark that mutual funds have achieved in a stock, averaged over time and across different mutual funds. The sample period runs from August 2007 to December 2023.

Statistics	Count	Mean	Std.Dev.	Min	5th pct.	10th pct.	25th pct.	Median	75th pct.	90th pct.	95th pct.	Max	corr(DFB,ADFD)
Panel I: S&P500 benchmark index													
			A. Disaggregate weight of stock holding in mutual fund portfolio (by month, mutual fund, and stock)										
Full sample	12696772	0.848	1.025	-4.350	0.030	0.050	0.140	0.480	1.190	2.120	2.870	54.290	-
Subsample DFB>0	9353057	1.064	1.085	0.010	0.080	0.130	0.290	0.740	1.480	2.430	3.150	54.290	-
Subsample DFB<0	3343715	0.242	0.441	-4.350	0.010	0.020	0.040	0.090	0.240	0.660	1.020	7.370	-
			B. Disaggregate weight of stock holding in benchmark index (by month, mutual fund, and stock)										
Full sample	12696772	0.314	0.511	0.003	0.031	0.041	0.066	0.129	0.330	0.852	1.259	7.372	-
Subsample DFB>0	9353057	0.273	0.428	0.003	0.030	0.040	0.063	0.120	0.288	0.702	1.079	7.372	-
Subsample DFB<0	3343715	0.428	0.679	0.003	0.034	0.045	0.076	0.165	0.473	1.164	1.632	7.372	-
			C. Disaggregate Deviation from Benchmark (by month, mutual fund, and stock) - DFB										
Full sample	12696772	0.534	0.911	-7.362	-0.258	-0.092	-0.004	0.226	0.852	1.664	2.288	51.828	0.886
Subsample DFB>0	9353057	0.792	0.907	0.000	0.019	0.043	0.151	0.491	1.115	1.932	2.574	51.828	1.000
Subsample DFB<0	3343715	-0.187	0.382	-7.362	-0.827	-0.503	-0.180	-0.055	-0.019	-0.007	-0.004	-0.000	-1.000
			D. Disaggregate Absolute Deviation from Benchmark (by month, mutual fund, and stock) - ADFB										
Full sample	12696772	0.632	0.846	0.000	0.008	0.018	0.068	0.301	0.892	1.688	2.315	51.828	-
Subsample DFB>0	9353057	0.792	0.907	0.000	0.019	0.043	0.151	0.491	1.115	1.932	2.574	51.828	-
Subsample DFB<0	3343715	0.187	0.382	0.000	0.004	0.007	0.019	0.055	0.180	0.503	0.827	7.362	-
Panel II: S&P400 benchmark index													
			A. Disaggregate weight of stock holding in mutual fund portfolio (by month, mutual fund, and stock)										
Full.sample	5816383	0.472	0.646	-5.680	0.010	0.020	0.050	0.210	0.650	1.270	1.750	20.820	-
Subsample DFB>0	2699040	0.905	0.733	0.010	0.190	0.250	0.400	0.710	1.170	1.810	2.300	20.820	-
Subsample DFB<0	3117343	0.098	0.105	-5.680	0.010	0.010	0.020	0.060	0.140	0.230	0.300	1.660	-
			B. Disaggregate weight of stock holding in benchmark index (by month, mutual fund, and stock)										
Full sample	5816383	0.281	0.139	0.001	0.107	0.130	0.180	0.255	0.352	0.468	0.549	1.702	-
Subsample DFB>0	2699040	0.265	0.132	0.001	0.101	0.123	0.169	0.240	0.331	0.442	0.520	1.702	-
Subsample DFB<0	3117343	0.295	0.143	0.001	0.113	0.137	0.192	0.270	0.369	0.488	0.569	1.702	-
			C. Disaggregate Deviation from Benchmark (by month, mutual fund, and stock) - DFB										
Full sample	5816383	0.191	0.644	-5.922	-0.388	-0.308	-0.186	-0.028	0.371	0.979	1.446	20.371	0.902
Subsample DFB>0	2699040	0.640	0.704	0.000	0.020	0.046	0.150	0.422	0.883	1.500	1.990	20.371	1.000
Subsample DFB<0	3117343	-0.198	0.137	-5.922	-0.456	-0.380	-0.273	-0.174	-0.095	-0.042	-0.021	-0.000	-1.000
			D. Disaggregate Absolute Deviation from Benchmark (by month, mutual fund, and stock) - ADFB										
Full sample	5816383	0.403	0.537	0.000	0.021	0.043	0.109	0.226	0.453	0.981	1.446	20.371	-
Subsample DFB>0	2699040	0.640	0.704	0.000	0.020	0.046	0.150	0.422	0.883	1.500	1.990	20.371	-
Subsample DFB<0	3117343	0.198	0.137	0.000	0.021	0.042	0.095	0.174	0.273	0.380	0.456	5.922	-

Table 2: **Descriptive statistics of stock active shares**

For each benchmark index (S&P500 in Panel I and S&P400 in Panel II), the table displays descriptive statistics of stock active shares (NSAS and SAS), considering all constituents with a mutual fund holding disclosure, and across all time periods. Reported statistics are the mean, standard deviation, minimum, percentiles, maximum, and correlation between NSAS and SAS.Each statistic is a single value that describes the typical active share of a stock, averaged over time. The sample period runs from August 2007 to December 2023.

Statistics	Count	Mean	Std.Dev.	Min	5th pct.	10th pct.	25th pct.	Median	75th pct.	90th pct.	95th pct.	Max	corr(NSAS,SAS)
Panel I: S&P500 benchmark index													
A. Aggregate weight of stock holdings in mutual fund portfolios (by month and stock)													
Full sample	97546	110.342	174.329	0.000	3.960	8.190	23.570	63.175	127.968	243.845	374.493	6010.120	-
Subsample NSAS>0	96680	107.264	154.947	0.010	4.100	8.300	23.598	62.870	126.760	238.212	362.652	4585.700	-
Subsample NSAS<0	866	453.887	790.251	0.000	0.020	0.060	12.012	221.780	547.790	968.170	1502.015	6010.120	-
B. Aggregate weight of stock holding in benchmark index (by month and stock)													
Full sample	97546	40.837	137.272	0.007	0.410	0.982	3.252	9.242	25.945	81.534	178.674	6752.325	-
Subsample NSAS>0	96680	36.524	100.350	0.007	0.423	0.995	3.249	9.171	25.392	76.803	160.998	4536.427	-
Subsample NSAS<0	866	522.446	874.725	0.009	0.050	0.095	14.783	255.862	643.725	1175.080	1898.332	6752.325	-
C. Aggregate Deviation from Benchmark (by month and stock) - NSAS													
Full sample	97546	69.504	76.324	-781.690	2.588	5.861	17.583	49.259	96.945	159.114	208.568	997.837	0.780
Subsample NSAS>0	96680	70.741	74.715	0.000	3.135	6.377	18.208	49.969	97.550	159.653	209.283	997.837	0.938
Subsample NSAS<0	866	-68.559	117.211	-781.690	-322.992	-206.094	-86.177	-16.834	-0.790	-0.026	-0.009	-0.000	-0.745
D. Aggregate Absolute Deviation from Benchmark (by month and stock) - SAS													
Full sample	97546	82.305	101.831	0.000	3.356	6.894	19.852	54.042	106.987	185.961	257.127	2898.011	-
Subsample NSAS>0	96680	82.183	102.562	0.000	3.334	6.854	19.871	54.069	106.778	184.845	254.669	2898.011	-
Subsample NSAS<0	866	83.107	96.888	0.001	3.494	7.166	19.732	53.835	108.570	192.329	273.444	1300.080	-
Panel II: S&P400 benchmark index													
A. Aggregate weight of stock holdings in mutual fund portfolios (by month and stock)													
Full sample	79349	34.601	31.883	0.000	1.870	3.700	9.780	25.870	50.160	77.232	97.210	296.580	-
Subsample NSAS>0	66738	37.399	32.462	0.030	2.320	4.510	12.060	29.530	53.470	80.700	100.551	296.580	-
Subsample NSAS<0	12611	19.794	23.634	0.000	0.460	1.710	4.700	11.170	25.710	51.050	70.025	236.230	-
B. Aggregate weight of stock holding in benchmark index (by month and stock)													
Full sample	79349	20.595	21.869	0.021	0.742	1.683	5.473	13.936	28.182	47.656	63.594	369.237	-
Subsample NSAS>0	66738	19.614	20.282	0.029	0.746	1.597	5.169	13.487	27.277	45.402	59.627	239.366	-
Subsample NSAS<0	12611	25.787	28.293	0.021	0.703	2.532	7.048	16.474	34.303	62.083	83.720	369.237	-
C. Aggregate Deviation from Benchmark (by month and stock) - NSAS													
Full sample	79349	14.006	18.268	-141.617	-6.556	-2.037	1.662	9.406	22.563	37.938	49.248	162.099	0.685
Subsample NSAS>0	66738	17.785	17.188	0.000	0.710	1.515	4.726	12.875	25.625	40.788	51.812	162.099	0.825
Subsample NSAS<0	12611	-5.993	7.796	-141.617	-20.328	-14.901	-8.087	-3.365	-1.080	-0.263	-0.112	-0.000	-0.611
D. Aggregate Absolute Deviation from Benchmark (by month and stock) - SAS													
Full sample	79349	29.524	26.300	0.000	1.470	3.124	8.647	22.663	43.168	65.324	81.224	220.468	-
Subsample NSAS>0	66738	29.661	26.342	0.000	1.510	3.184	8.821	22.792	43.328	65.533	81.596	220.468	-
Subsample NSAS<0	12611	29.126	26.174	0.000	1.370	2.957	8.185	22.317	42.740	64.673	80.190	214.320	-

Table 3: **Performance of the SAS-Oracle strategy - with quarterly rebalancing.**
This table shows the performance of the SAS-Oracle strategy implemented by using the real value of stocks SAS values as computed from observed data. The strategy consists of using the computed market-wide absolute deviation from the benchmark for each stock at time t as if it was known, then sorting the stocks into quantile portfolios based on their computed SAS values. The formed portfolio is held for three months and rebalanced in the beginning of each quarter. *Oracle* portfolios are formed using ex-ante the observed (ex-post) measure of stocks' SAS computed using mutual fund holdings when released to sort stocks into quantile portfolios. The evaluation period goes from October 2007 to December 2023.

	Decile portfolios										Quintile portfolios					Tercile portfolios			Benchmark
Statistics	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q1	Q2	Q3	Q4	Q5	Q1	Q2	Q3	Index
Panel I: S&P500 benchmark index																			
Mean	0.817	1.110	1.233	1.426	1.146	1.407	1.288	1.567	1.548	1.561	0.986	1.345	1.283	1.435	1.560	1.151	1.310	1.530	0.864
Std.Dev.	7.102	6.167	5.931	5.375	5.094	5.215	5.188	4.888	4.637	4.243	6.474	5.509	5.087	4.940	4.351	6.066	5.094	4.513	4.662
Sharpe.Ratio	0.104	0.168	0.195	0.251	0.210	0.255	0.233	0.305	0.317	0.350	0.140	0.230	0.237	0.275	0.341	0.177	0.242	0.322	0.169
Skew	0.102	-0.116	0.068	-0.279	-0.476	-0.404	-0.370	-0.103	-0.196	-0.228	-0.059	-0.193	-0.455	-0.259	-0.227	-0.087	-0.396	-0.236	-0.554
Kurtosis	3.448	3.092	3.376	1.745	1.536	1.739	1.482	0.579	0.599	0.324	3.018	2.270	1.665	0.721	0.485	2.999	1.588	0.509	0.810
Adjusted.Sharpe.Ratio	0.104	0.166	0.194	0.247	0.206	0.249	0.229	0.303	0.313	0.345	0.140	0.227	0.232	0.271	0.336	0.176	0.237	0.317	0.166
Alpha	-0.133	0.221	0.274	0.547	0.254	0.495	0.408	0.671	0.697	0.744	0.069	0.428	0.382	0.547	0.728	0.231	0.413	0.675	-0.018
s.e.alpha.	0.187	0.128	0.116	0.151	0.094	0.100	0.091	0.108	0.105	0.117	0.128	0.113	0.072	0.085	0.093	0.113	0.075	0.081	0.017
Panel II: S&P400 benchmark index																			
Mean	1.270	1.374	1.341	1.306	1.370	1.547	1.385	1.485	1.802	2.363	1.354	1.330	1.469	1.441	2.107	1.366	1.425	1.871	0.864
Std.Dev.	8.286	7.024	6.321	6.236	6.243	5.952	5.904	5.593	5.633	5.428	7.381	6.155	5.967	5.663	5.407	6.810	5.884	5.407	5.886
Sharpe.Ratio	0.144	0.185	0.200	0.197	0.207	0.247	0.222	0.252	0.306	0.421	0.173	0.204	0.233	0.241	0.375	0.189	0.229	0.332	0.134
Skew	1.030	0.119	0.065	-0.253	0.194	-0.325	-0.342	-0.317	-0.268	-0.067	0.475	-0.196	-0.133	-0.331	-0.157	0.252	-0.191	-0.235	-0.567
Kurtosis	7.692	3.279	2.516	1.253	2.910	1.415	1.558	0.551	0.792	0.318	4.778	1.422	1.856	0.965	0.530	3.400	1.610	0.560	2.033
Adjusted.Sharpe.Ratio	0.147	0.185	0.200	0.195	0.208	0.243	0.218	0.248	0.301	0.418	0.174	0.202	0.231	0.237	0.370	0.190	0.227	0.327	0.132
Alpha	0.270	0.380	0.381	0.354	0.441	0.629	0.446	0.569	0.838	1.413	0.358	0.374	0.546	0.514	1.150	0.382	0.502	0.924	-0.067
s.e.alpha.	0.195	0.143	0.151	0.139	0.148	0.138	0.116	0.136	0.180	0.203	0.119	0.116	0.114	0.114	0.178	0.095	0.096	0.148	0.095

Table 4: **Details of characteristics**

This table lists the explanatory variables used in our machine learning models to predict stock selection into SAS portfolio. The variables are inspired by Green et al. (2017) and Gu et al. (2020).

No.	Acronym	Firm Characteristic	Paper's author(s)	Year, Journal	Data Source	Frequency
Panel A: Stock predictors						
1	absacc	Absolute accruals	Bandyopadhyay, Huang& Wirjanto	2010, WP	Compustat	Annual
2	acc	Working capital accruals	Sloan	1996, TAR	Compustat	Annual
3	age	# years since first Compustat coverage	Jiang, Lee & Zhang	2005, RAS	Compustat	Annual
4	agr	Asset growth	Cooper, Gulen & Schill	2008, JF	Compustat	Annual
5	beta12m	12-month rolling market risk's exposure	Bacon	2009, WP	CRSP	Monthly
6	beta24m	24-month rolling market risk's exposure	Bacon	2009, WP	CRSP	Monthly
7	beta36m	36-month rolling market risk's exposure	Bacon	2009, WP	CRSP	Monthly
8	bms	Book-to-market	Rosenberg, Reid & Lanstein	1985, JPM	Compustat+CRSP	Annual
9	bm_ia	Industry-adjusted book to market	Asness, Porter & Stevens	2000, WP	Compustat+CRSP	Annual
10	cash	Cash holdings	Palazzo	2012, JFE	Compustat	Quarterly
11	cashdebt	Cash flow to debt	Ou & Penman	1989, JAE	Compustat	Annual
12	cashpr	Cash productivity	Chandrashekar & Rao	2009, WP	Compustat	Annual
13	cfp	Cash flow to price ratio	Desai, Rajgopal & Venkatachalam	2004, TAR	Compustat	Annual
14	cfp_ia	Industry-adjusted cash flow to price ratio	Asness, Porter & Stevens	2000, WP	Compustat	Annual
15	chatoia	Industry-adjusted change in asset turnover	Soliman	2008, TAR	Compustat	Annual
16	chcsho	Change in shares outstanding	Pontiff & Woodgate	2008, JF	Compustat	Annual
17	chempia	Industry-adjusted change in employees	Asness, Porter & Stevens	1994, WP	Compustat	Annual
18	chinv	Change in inventory	Thomas & Zhang	2002, RAS	Compustat	Annual
19	chmom	Change in 6-month momentum	Gettleman & Marks	2006, WP	CRSP	Monthly
20	chpmia	Industry-adjusted change in profit margin	Soliman	2008, TAR	Compustat	Annual
21	chtx	Change in tax expense	Thomas & Zhang	2011, JAR	Compustat	Quarterly
22	cinvest	Corporate investment	Titman, Wei & Xie	2004, JFQA	Compustat	Quarterly
23	convind	Convertible debt indicator	Valta	2016, JFQA	Compustat	Annual
24	currat	Current ratio	Ou & Penman	1989, JAE	Compustat	Annual
25	depr	Depreciation / PP&E	Holthausen & Larcker	1992, JAE	Compustat	Annual
26	divi	Dividend initiation	Michaely, Thaler & Womack	1995, JF	Compustat	Annual
27	divo	Dividend omission	Michaely, Thaler & Womack	1995, JF	Compustat	Annual
28	dy	Dividend to price	Litzenberger & Ramaswamy	1982, JF	Compustat	Annual
29	egr	Growth in common shareholder equity	Richardson, Sloan, Soliman & Tuna	2005, JAE	Compustat	Annual
30	ep	Earnings to price	Basu	1977, JF	Compustat	Annual
31	gma	Gross profitability	Novy-Marx	2013, JFE	Compustat	Annual
32	grCAPX	Growth in capital expenditures	Anderson & Garcia-Feijoo	2006, JF	Compustat	Annual
33	herf	Industry sales concentration	Hou & Robinson	2006, JF	Compustat	Annual
34	hire	Employee growth rate	Bazdresch, Belo & Lin	2014, JPE	Compustat	Annual
35	invest	Capital expenditures and inventory	Chen & Zhang	2010, JF	Compustat	Annual
36	ir12m	12-month rolling information Ratio	Bacon	2009, WP	CRSP	Monthly
37	ir24m	24-month rolling information Ratio	Bacon	2009, WP	CRSP	Monthly
38	ir36m	36-month rolling information Ratio	Bacon	2009, WP	CRSP	Monthly
39	lev	Leverage	Bhandari	1988, JF	Compustat	Annual
40	lgr	Growth in long-term debt	Richardson, Sloan, Soliman & Tuna	2005, JAE	Compustat	Annual
41	maxret	Maximum daily return	Bali, Cakici & Whitelaw	2011, JFE	CRSP	Monthly
42	mc	Market value of common equity (csho*prcc_f)	Banz	1981, JFE	CRSP	Monthly
43	mom12m	12-month momentum	Jegadeesh	1990, JF	CRSP	Monthly
44	mom1m	1-month momentum	Jegadeesh & Titman	1993, JF	CRSP	Monthly
45	mom36m	36-month momentum	Jegadeesh & Titman	1993, JF	CRSP	Monthly
46	mom6m	6-month momentum	Jegadeesh & Titman	1993, JF	CRSP	Monthly
47	ms	Financial statement score	Mohanram	2005, RAS	Compustat	Quarterly
48	mve	Logarithm of firm size	Banz	1981, JFE	CRSP	Monthly
49	mve_ia	Industry-adjusted size	Asness, Porter & Stevens	2000, WP	Compustat	Annual
50	nincr	Number of earnings increases	Barth, Elliott & Finn	1999, JAR	Compustat	Quarterly
51	operprof	Operating profitability	Fama & French	2015, JFE	Compustat	Annual
52	orgcap	Organizational capital	Eisfeldt & Papanikolaou	2013, JF	Compustat	Annual
53	pchcapx_ia	Industry adjusted % change in capital expenditures	Abarbanell & Bushee	1998, TAR	Compustat	Annual
54	pchcurrat	% change in current ratio	Ou & Penman	1989, JAE	Compustat	Annual
55	pchdepr	% change in depreciation	Holthausen & Larcker	1992, JAE	Compustat	Annual
56	pchgm_pchsale	% change in gross margin - % change in sales	Abarbanell & Bushee	1998, TAR	Compustat	Annual
57	pchquick	% change in quick ratio	Ou & Penman	1989, JAE	Compustat	Annual
58	pchsale_pchinvt	% change in sales - % change in inventory	Abarbanell & Bushee	1998, TAR	Compustat	Annual
59	pchsale_pchrect	% change in sales - % change in A/R	Abarbanell & Bushee	1998, TAR	Compustat	Annual
60	pchsale_pchxsga	% change in sales - % change in SG&A	Abarbanell & Bushee	1998, TAR	Compustat	Annual
61	pchsaleinv	% change sales-to-inventory	Ou & Penman	1989, JAE	Compustat	Annual

Table 4: **Details of characteristics (continued)**
This table lists the explanatory variables used in our machine learning models to predict stock selection into SAS portfolio. The variables are inspired by Green et al. (2017) and Gu et al. (2020).

No.	Acronym	Firm Characteristic	Paper's author(s)	Year, Journal	Data Source	Frequency
62	pctacc	Percent accruals	Hafzalla, Lundholm & Van Winkle	2011, TAR	Compustat	Annual
63	quick	Quick ratio	Ou & Penman	1989, JAE	Compustat	Annual
64	rd	R&D increase	Eberhart, Maxwell & Siddique	2004, JF	Compustat	Annual
65	rd_mve	R&D to market capitalization	Guo, Lev & Shi	2006, JBFA	Compustat	Annual
66	rd_sale	R&D to sales	Guo, Lev & Shi	2006, JBFA	Compustat	Annual
67	realestate	Real estate holdings	Tuzel	2010, RFS	Compustat	Annual
68	roaq	Return on assets	Balakrishnan, Bartov & Faurel	2010, JAE	Compustat	Quarterly
69	roavol	Earnings volatility	Francis, LaFond, Olsson & Schipper	2004, TAR	Compustat	Quarterly
70	roeq	Return on equity	Hou, Xue & Zhang	2015, RFS	Compustat	Quarterly
71	roic	Return on invested capital	Brown & Rowe	2007, WP	Compustat	Annual
72	rsup	Revenue surprise	Kama	2009, JBFA	Compustat	Quarterly
73	salecash	Sales to cash	Ou & Penman	1989, JAE	Compustat	Annual
74	saleinv	Sales to inventory	Ou & Penman	1989, JAE	Compustat	Annual
75	salerec	Sales to receivables	Ou & Penman	1989, JAE	Compustat	Annual
76	secured	Secured debt	Valta	2016, JFQA	Compustat	Annual
77	securedind	Secured debt indicator	Valta	2016, JFQA	Compustat	Annual
78	sgr	Sales growth	Lakonishok, Shleifer & Vishny	1994, JF	Compustat	Annual
79	sic_2	2-digit SIC code	Green, Hand, & Zhang	2017, RFS	Compustat	Annual
80	sin	Sin stocks	Hong & Kacperczyk	2009, JFE	Compustat	Annual
81	sp	Sales to price	Barbee, Mukherji, & Raines	1996, FAJ	Compustat	Annual
82	sr12m	12-month rolling Sharpe Ratio	Bacon	2009, WP	CRSP	Monthly
83	sr24m	24-month rolling Sharpe Ratio	Bacon	2009, WP	CRSP	Monthly
84	sr36m	36-month rolling Sharpe Ratio	Bacon	2009, WP	CRSP	Monthly
85	stdacc	Accrual volatility	Bandyopadhyay, Huang & Wirjanto	2010, WP	Compustat	Quarterly
86	stdcf	Cash flow volatility	Huang	2009, JEF	Compustat	Quarterly
87	tang	Debt capacity/firm tangibility	Almeida & Campello	2007, RFS	Compustat	Annual
88	tb	Tax income to book income	Lev & Nissim	2004, TAR	Compustat	Annual
89	tr12m	12-month rolling Treynor Ratio	Bacon	2009, WP	CRSP	Monthly
90	tr24m	24-month rolling Treynor Ratio	Bacon	2009, WP	CRSP	Monthly
91	tr36m	36-month rolling Treynor Ratio	Bacon	2009, WP	CRSP	Monthly
Panel B: Macroeconomic predictors						
92	D12	Log dividend to price ratio	Goyal & Welch	2007, RFS	CRSP	Monthly
93	E12	Log earning to price ratio	Goyal & Welch	2007, RFS	CRSP	Monthly
94	bm	Book value to market value for the DJIA	Goyal & Welch	2007, RFS	CRSP	Monthly
95	dfy	Default Yield Spread	Goyal & Welch	2007, RFS	CRSP	Monthly
96	ntis	Net Equity Expansion	Goyal & Welch	2007, RFS	CRSP	Monthly
97	tbl	Treasury-bill rates	Goyal & Welch	2007, RFS	CRSP	Monthly
98	tms	Term Spread	Goyal & Welch	2007, RFS	CRSP	Monthly
99	svar	Stock Variance	Goyal & Welch	2007, RFS	CRSP	Monthly

Table 5: **Out-of-sample measures of prediction performance**

This table shows goodness-of-fit summary statistics of the machine learning models. The dependent variable is the logarithm of SAS value, and predictors are macrovariables and stock characteristics summarized in Table 4. We consider using each benchmark index separately. The following summary statistics are displayed: Quantile (decile, quintile, tercile) portfolio number, coefficient of variation of the dependent variable (CV) modified out-of-sample r-squared ($R^2_{OOS_M}$) defined in equation (5), traditional out-of-sample r-squared ($R^2_{OOS_T}$) defined in equation (4). Quantile portfolios are formed by sorting stocks based on the predicted variable. The initial training period is from October 2007 to December 2017. The training is updated recursively every quarter by expanding the training period, and the testing period is a quarter ahead in the future. All the statistics are computed over the out-of-sample prediction period from January 2018 to Dec. 2023.

	Random Forest					Gradient Boosting					OLS					Elastic Net					Deep Neural Network				
	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$
Panel I: S&P500 benchmark index																									
I.A. Decile portfolios																									
1	3.33	3.28	0.24	0.13	0.12	3.29	3.37	0.24	-0.04	-0.06	3.37	3.73	0.24	-0.09	-0.29	3.37	3.74	0.24	-0.08	-0.29	3.32	3.47	0.25	0.03	-0.01
2	3.71	3.60	0.21	0.20	0.18	3.70	3.80	0.20	0.13	0.11	3.68	3.88	0.20	0.05	-0.03	3.68	3.88	0.20	0.06	-0.01	3.71	3.82	0.20	0.00	-0.02
3	3.88	3.76	0.19	0.24	0.21	3.86	3.96	0.19	0.12	0.10	3.85	3.95	0.19	0.07	0.05	3.85	3.95	0.19	0.09	0.07	3.85	3.98	0.19	0.09	0.06
4	3.98	3.88	0.19	0.25	0.23	3.98	4.08	0.19	0.14	0.13	3.98	4.01	0.19	0.08	0.08	3.98	4.02	0.19	0.09	0.09	3.98	4.10	0.19	0.10	0.08
5	4.11	3.99	0.19	0.27	0.25	4.11	4.17	0.18	0.16	0.15	4.09	4.07	0.19	0.12	0.12	4.09	4.07	0.19	0.13	0.13	4.11	4.20	0.18	0.12	0.10
6	4.22	4.10	0.18	0.25	0.22	4.22	4.27	0.18	0.13	0.12	4.20	4.14	0.18	0.15	0.14	4.21	4.14	0.18	0.15	0.14	4.22	4.30	0.18	0.15	0.14
7	4.35	4.22	0.18	0.26	0.23	4.35	4.39	0.18	0.18	0.17	4.35	4.23	0.18	0.15	0.12	4.35	4.22	0.18	0.16	0.13	4.35	4.41	0.18	0.19	0.18
8	4.53	4.36	0.18	0.27	0.23	4.54	4.54	0.17	0.18	0.18	4.54	4.35	0.18	0.18	0.12	4.54	4.34	0.18	0.19	0.13	4.52	4.56	0.18	0.23	0.23
9	4.81	4.56	0.17	0.26	0.16	4.82	4.74	0.16	0.21	0.20	4.80	4.55	0.16	0.20	0.10	4.80	4.55	0.16	0.21	0.10	4.82	4.76	0.16	0.25	0.25
10	5.37	4.94	0.16	0.34	0.07	5.42	5.18	0.15	0.30	0.22	5.42	5.36	0.15	-0.49	-0.49	5.42	5.36	0.15	-0.50	-0.50	5.41	5.10	0.16	0.36	0.23
I.B. Quintile portfolios																									
1	3.52	3.44	0.23	0.20	0.19	3.50	3.59	0.23	0.10	0.09	3.52	3.80	0.23	-0.00	-0.13	3.52	3.81	0.23	0.01	-0.12	3.51	3.65	0.23	0.07	0.04
2	3.93	3.82	0.19	0.25	0.22	3.92	4.02	0.19	0.14	0.12	3.92	3.98	0.19	0.08	0.07	3.91	3.98	0.19	0.09	0.08	3.92	4.04	0.19	0.10	0.08
3	4.16	4.04	0.18	0.27	0.24	4.16	4.22	0.18	0.15	0.14	4.15	4.11	0.19	0.14	0.14	4.15	4.11	0.19	0.14	0.14	4.16	4.25	0.18	0.14	0.13
4	4.44	4.29	0.18	0.28	0.24	4.44	4.46	0.18	0.19	0.19	4.44	4.28	0.18	0.17	0.14	4.44	4.28	0.18	0.18	0.15	4.43	4.48	0.18	0.22	0.22
5	5.09	4.75	0.17	0.36	0.21	5.11	4.96	0.17	0.34	0.31	5.11	4.95	0.17	-0.02	-0.05	5.11	4.95	0.17	-0.02	-0.06	5.11	4.93	0.17	0.37	0.33
I.C. Tercile portfolios																									
1	3.67	3.57	0.22	0.25	0.24	3.65	3.74	0.22	0.16	0.14	3.66	3.87	0.22	0.06	-0.01	3.66	3.87	0.22	0.07	-0.00	3.65	3.79	0.22	0.12	0.10
2	4.16	4.04	0.19	0.27	0.25	4.16	4.22	0.18	0.17	0.16	4.15	4.11	0.19	0.14	0.14	4.15	4.11	0.19	0.15	0.15	4.16	4.25	0.18	0.16	0.14
3	4.85	4.58	0.18	0.39	0.30	4.87	4.78	0.18	0.36	0.35	4.86	4.70	0.18	0.16	0.12	4.87	4.70	0.18	0.16	0.13	4.86	4.77	0.18	0.38	0.37
All	4.23	4.07	0.23	0.48	0.45	4.23	4.25	0.23	0.44	0.44	4.23	4.23	0.23	0.33	0.33	4.23	4.23	0.23	0.33	0.33	4.23	4.27	0.23	0.43	0.43
Panel II: S&P400 benchmark index																									
II.A. Decile portfolios																									
1	2.70	2.25	0.32	-0.07	-0.35	2.72	2.30	0.32	-0.17	-0.40	2.65	2.80	0.32	-0.06	-0.08	2.64	2.81	0.33	-0.03	-0.07	2.68	2.53	0.32	-0.12	-0.14
2	2.97	2.53	0.28	-0.01	-0.30	2.97	2.61	0.28	-0.16	-0.35	2.88	2.96	0.28	-0.00	-0.01	2.88	2.96	0.28	0.03	0.02	2.86	2.89	0.28	-0.14	-0.15
3	3.13	2.68	0.27	0.02	-0.27	3.10	2.76	0.27	-0.08	-0.24	3.05	3.06	0.26	0.02	0.02	3.04	3.06	0.26	0.06	0.06	3.07	3.10	0.26	-0.10	-0.10
4	3.24	2.79	0.26	0.05	-0.25	3.19	2.87	0.25	-0.09	-0.25	3.18	3.15	0.25	0.05	0.05	3.18	3.16	0.25	0.07	0.06	3.18	3.25	0.25	-0.07	-0.07
5	3.36	2.88	0.25	0.07	-0.26	3.32	2.98	0.24	-0.10	-0.28	3.30	3.25	0.24	0.04	0.03	3.30	3.25	0.24	0.07	0.07	3.32	3.35	0.23	-0.03	-0.03
6	3.46	2.97	0.23	0.11	-0.26	3.40	3.08	0.24	-0.06	-0.22	3.41	3.34	0.23	0.07	0.06	3.41	3.34	0.23	0.08	0.07	3.41	3.44	0.23	-0.02	-0.02
7	3.53	3.04	0.23	0.10	-0.26	3.51	3.17	0.23	-0.10	-0.26	3.53	3.45	0.22	0.03	0.02	3.53	3.44	0.22	0.08	0.07	3.53	3.54	0.21	0.03	0.03
8	3.59	3.12	0.23	0.11	-0.22	3.63	3.27	0.22	-0.09	-0.30	3.67	3.58	0.21	0.02	0.00	3.67	3.58	0.20	0.06	0.05	3.67	3.63	0.21	0.03	0.03
9	3.68	3.21	0.21	0.10	-0.26	3.75	3.37	0.21	-0.08	-0.32	3.82	3.75	0.19	-0.06	-0.07	3.82	3.75	0.19	-0.01	-0.02	3.79	3.74	0.21	0.08	0.08
10	3.79	3.39	0.20	0.07	-0.22	3.88	3.57	0.19	-0.10	-0.27	3.99	4.14	0.19	-0.17	-0.21	3.99	4.13	0.20	-0.13	-0.16	3.95	3.90	0.20	0.13	0.13

Table 5: **Out-of-sample measures of prediction performance (continued)**

This table shows goodness-of-fit summary statistics of the machine learning models. The dependent variable is the logarithm of SAS value, and predictors are macrovariables and stock characteristics summarized in Table 4. We consider using each benchmark index separately. The following summary statistics are displayed: Quantile (decile, quintile, tercile) portfolio number, coefficient of variation of the dependent variable (CV), modified out-of-sample r-squared ($R^2_{OOS_M}$) defined in equation (5), traditional out-of-sample r-squared ($R^2_{OOS_T}$) defined in equation (4). Quantile portfolios are formed by sorting stocks based on the predicted variable. The initial training period is from October 2007 to December 2017. The training is updated recursively every quarter by expanding the training period, and the testing period is a quarter ahead in the future. All the statistics are computed over the out-of-sample prediction period from January 2018 to Dec. 2023.

	Random Forest					Gradient Boosting					OLS					Elastic Net					Deep Neural Network				
	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$	DepVar	PredVar	CV	$R^2_{OOS_M}$	$R^2_{OOS_T}$
Panel II : S&P400 benchmark index																									
II.B. Quintile portfolios																									
1	2.84	2.39	0.30	-0.02	-0.29	2.84	2.45	0.30	-0.15	-0.35	2.76	2.88	0.30	-0.01	-0.03	2.76	2.88	0.30	0.01	-0.01	2.77	2.71	0.30	-0.13	-0.13
2	3.18	2.73	0.26	0.04	-0.25	3.14	2.82	0.26	-0.08	-0.24	3.11	3.11	0.26	0.04	0.04	3.11	3.11	0.26	0.07	0.07	3.13	3.17	0.26	-0.08	-0.08
3	3.41	2.93	0.24	0.09	-0.25	3.36	3.03	0.24	-0.08	-0.25	3.35	3.29	0.23	0.06	0.05	3.36	3.29	0.23	0.08	0.07	3.36	3.40	0.23	-0.02	-0.02
4	3.56	3.08	0.23	0.10	-0.24	3.57	3.22	0.23	-0.08	-0.27	3.60	3.51	0.21	0.03	0.02	3.60	3.51	0.21	0.08	0.06	3.60	3.58	0.21	0.04	0.04
5	3.74	3.30	0.21	0.09	-0.24	3.82	3.47	0.20	-0.09	-0.29	3.90	3.94	0.19	-0.12	-0.13	3.90	3.94	0.20	-0.08	-0.08	3.87	3.82	0.20	0.11	0.11
II.C. Tercile portfolios																									
1	2.96	2.51	0.29	0.03	-0.24	2.95	2.58	0.29	-0.09	-0.27	2.89	2.96	0.29	0.03	0.02	2.88	2.96	0.29	0.06	0.05	2.90	2.87	0.29	-0.08	-0.08
2	3.40	2.92	0.24	0.10	-0.23	3.35	3.03	0.24	-0.08	-0.24	3.35	3.29	0.24	0.07	0.06	3.35	3.29	0.23	0.09	0.08	3.36	3.39	0.23	0.00	-0.00
3	3.68	3.22	0.22	0.10	-0.23	3.73	3.38	0.21	-0.07	-0.27	3.80	3.79	0.20	-0.04	-0.04	3.79	3.78	0.20	0.00	0.00	3.78	3.74	0.21	0.10	0.10
All	3.34	2.88	0.26	0.18	-0.10	3.34	3.00	0.26	0.06	-0.10	3.34	3.34	0.26	0.19	0.19	3.34	3.34	0.26	0.22	0.22	3.34	3.33	0.26	0.17	0.17

Table 6: **Investment Performance of Machine Learning Algorithms - with quarterly rebalancing.**

This table shows the performance of the machine learning based SAS investment strategy. The Strategy consists of predicting stock future SAS value based on its characteristics at time t, then sorting the stocks into quantile portfolios, and taking a long position in top quantile portfolios composed of stocks with the highest predicted SAS values. The prediction methods are random forest (RF), gradient boosting (GB), ordinary least squares (OLS), elastic net (ENET), deep neural network (DNN). We also use analyst recommendation consensus (AR) to form the portfolio in a similar fashion to the SAS strategy, and the results are presented in the AR columns. ORA columns show the statistics for the SAS-Oracle strategy implemented by using the real value of stocks SAS values as computed from observed data. The formed portfolio is held for three months, it is rebalanced in the beginning of the next quarter, and the investment runs from January 2018 to December 2023.

	Top 10							Top 5							Top 3							Benchmark
Statistics	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	Index
Panel I: S&P500 benchmark index																						
Mean	1.315	1.345	1.344	1.357	1.390	1.414	1.420	1.366	1.390	1.359	1.369	1.394	1.354	1.396	1.335	1.400	1.376	1.381	1.368	1.343	1.371	1.089
Std. Dev.	5.175	5.003	4.949	4.953	5.015	5.884	5.698	5.126	5.107	5.114	5.132	5.099	5.770	5.789	5.306	5.221	5.224	5.230	5.235	5.770	5.718	5.213
Sharpe Ratio	0.225	0.239	0.241	0.244	0.248	0.215	0.223	0.237	0.243	0.237	0.238	0.244	0.209	0.215	0.223	0.239	0.235	0.236	0.233	0.207	0.214	0.180
Skew	-0.091	-0.223	-0.181	-0.190	-0.177	-0.062	-0.061	-0.039	-0.116	-0.113	-0.113	-0.085	-0.272	-0.171	-0.068	-0.106	-0.112	-0.110	-0.116	-0.166	-0.199	-0.370
Kurtosis	-0.289	-0.092	-0.185	-0.192	-0.288	0.650	-0.204	-0.176	-0.191	-0.136	-0.169	-0.158	0.817	-0.094	-0.101	-0.117	-0.103	-0.093	-0.095	0.481	0.048	-0.206
Adjusted Sharpe Ratio	0.225	0.237	0.240	0.242	0.246	0.214	0.223	0.237	0.242	0.236	0.237	0.243	0.206	0.214	0.223	0.238	0.234	0.235	0.232	0.205	0.212	0.178
Alpha	0.222	0.312	0.277	0.291	0.348	0.309	0.438	0.296	0.332	0.291	0.298	0.346	0.280	0.398	0.263	0.351	0.322	0.324	0.322	0.241	0.368	0.004
s.e.(alpha)	0.098	0.069	0.081	0.082	0.072	0.145	0.124	0.089	0.079	0.083	0.085	0.090	0.121	0.108	0.103	0.095	0.086	0.088	0.084	0.110	0.116	0.028
Panel II: S&P400 benchmark index																						
Mean	1.605	1.552	1.807	1.768	1.739	1.279	1.967	1.644	1.666	1.766	1.745	1.700	1.402	1.819	1.626	1.652	1.695	1.700	1.674	1.442	1.641	0.906
Std.Dev.	6.377	6.022	6.310	6.311	6.120	6.629	6.023	6.471	6.292	6.399	6.398	6.280	6.748	6.118	6.397	6.413	6.481	6.499	6.402	6.825	6.089	6.838
Sharpe Ratio	0.228	0.233	0.262	0.256	0.260	0.170	0.302	0.231	0.241	0.253	0.249	0.247	0.186	0.273	0.231	0.234	0.238	0.238	0.238	0.189	0.245	0.111
Skew	-0.217	-0.173	-0.063	-0.048	-0.025	-0.302	-0.112	-0.179	-0.135	-0.075	-0.065	-0.021	-0.196	-0.126	-0.108	-0.103	0.012	0.043	-0.065	-0.136	-0.225	-0.515
Kurtosis	0.183	-0.341	-0.318	-0.292	-0.214	0.080	-0.360	-0.062	-0.251	-0.319	-0.331	-0.276	0.032	-0.319	0.045	-0.112	0.114	0.149	0.079	0.229	-0.234	1.377
Adjusted Sharpe Ratio	0.226	0.231	0.262	0.256	0.260	0.169	0.300	0.229	0.240	0.252	0.249	0.247	0.184	0.271	0.230	0.233	0.238	0.239	0.237	0.188	0.243	0.110
Alpha	0.471	0.551	0.709	0.683	0.684	0.178	0.887	0.511	0.600	0.663	0.642	0.616	0.335	0.731	0.522	0.595	0.590	0.588	0.578	0.345	0.584	-0.060
s.e.(alpha)	0.180	0.194	0.208	0.216	0.198	0.251	0.214	0.187	0.211	0.211	0.203	0.198	0.206	0.183	0.134	0.187	0.181	0.186	0.182	0.141	0.152	0.131

Table 7: **Additional Investment Performance Measures of Machine Learning Algorithms - with quarterly rebalancing.**
This table shows additional performance measures of the machine learning based investment strategy to account for downside risks. The Strategy consists of predicting stock future SAS value based on its characteristics at time t, then sorting the stocks into quantile portfolios, and investing in top quantile portfolios composed of stocks with the highest predicted SAS values. The prediction methods are random forest (RF), gradient boosting (GB), ordinary least squares (OLS), elastic net (ENET), deep neural network (DNN). We also use analyst recommendation consensus (AR) to form the portfolio in a similar fashion to the SAS strategy, and the results are presented in the AR columns. ORA columns show the statistics for the SAS-Oracle strategy implemented by using the real value of stocks SAS values as computed from observed data. The formed portfolio is held for three months, it is rebalanced in the beginning of the next quarter, and the investment runs from January 2018 to December 2023.

	Top 10							Top 5							Top 3							Benchmark
Statistics	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	Index
Panel I: S&P500 benchmark index																						
Sortino	0.373	0.391	0.396	0.400	0.411	0.351	0.377	0.402	0.409	0.393	0.396	0.411	0.330	0.354	0.374	0.403	0.393	0.394	0.389	0.332	0.349	0.275
Info Ratio	0.240	0.393	0.328	0.343	0.411	0.235	0.273	0.316	0.393	0.342	0.344	0.375	0.257	0.306	0.288	0.396	0.373	0.384	0.376	0.242	0.309	0.016
Drawdown	0.232	0.214	0.238	0.238	0.222	0.218	0.224	0.200	0.186	0.207	0.206	0.199	0.260	0.227	0.202	0.191	0.195	0.197	0.201	0.243	0.234	0.243
VaR	0.095	0.095	0.089	0.089	0.091	0.117	0.107	0.089	0.095	0.092	0.091	0.091	0.120	0.115	0.099	0.098	0.099	0.099	0.098	0.116	0.115	0.103
Turnover	0.474	0.437	0.433	0.433	0.435	0.481	0.489	0.422	0.412	0.407	0.406	0.402	0.430	0.456	0.377	0.368	0.365	0.364	0.362	0.370	0.387	0.462
Panel II: S&P400 benchmark index																						
Sortino	0.373	0.390	0.460	0.448	0.457	0.264	0.538	0.384	0.409	0.439	0.434	0.432	0.295	0.474	0.387	0.396	0.412	0.414	0.406	0.303	0.408	0.161
Info Ratio	0.263	0.282	0.362	0.356	0.375	0.089	0.405	0.314	0.339	0.369	0.362	0.360	0.207	0.396	0.375	0.371	0.387	0.382	0.383	0.258	0.382	-0.054
Drawdown	0.245	0.207	0.204	0.215	0.185	0.273	0.230	0.241	0.208	0.206	0.205	0.198	0.271	0.222	0.250	0.244	0.233	0.230	0.242	0.277	0.230	0.354
VaR	0.130	0.119	0.120	0.119	0.115	0.148	0.105	0.126	0.124	0.128	0.127	0.125	0.148	0.114	0.123	0.125	0.125	0.125	0.127	0.142	0.119	0.153
Turnover	0.625	0.619	0.581	0.582	0.581	0.539	0.564	0.553	0.554	0.523	0.524	0.527	0.478	0.509	0.479	0.473	0.446	0.446	0.449	0.423	0.431	0.382

Table 8: **Linear Panel Regression Models of SAS values due to Carbon Emissions Level and Intensity**

This table shows the regressions of logarithm of SAS measure (the dependent variable) on logarithm of carbon emission intensity (Panel A) or on logarithm of carbon emission level (Panel B). Carbon emission intensity is computed as the ratio of CO_2 emission (scope 1 & 2) to stock market capitalization. The panel data runs from 2016 to 2023. Standard errors are clustered at stock level. FE stands for fixed effects. *, **, and *** denote respectively a p-value lower than 0.1, 0.05, and 0.01.

Panel A: Carbon Intensity is the main explanatory variable

	S&P400			S&P500		
	(1)	(2)	(3)	(1)	(2)	(3)
Emission Intensity	-0.029^{***}	-0.110^{***}	-0.113^{***}	-0.052^{***}	-0.116^{***}	-0.142^{***}
	(0.005)	(0.014)	(0.015)	(0.005)	(0.020)	(0.023)
Constant	3.926^{***}	4.688^{***}	4.590^{***}	5.207^{***}	6.828^{***}	7.856^{***}
	(0.076)	(0.239)	(0.262)	(0.076)	(0.433)	(0.555)
Stock FE	NO	YES	YES	NO	YES	YES
Time FE	NO	NO	YES	NO	NO	YES
Observations	1,798	1,798	1,798	2,182	2,182	2,182
R^2	0.017	0.803	0.813	0.049	0.896	0.904
Adjusted R^2	0.016	0.714	0.727	0.048	0.866	0.875
Residual Std. Error	0.559 (df = 1796)	0.301 (df = 1235)	0.294 (df = 1228)	0.662 (df = 2180)	0.248 (df = 1691)	0.240 (df = 1682)
F Statistic	30.631^{***} (df = 1; 1796)	8.983^{***} (df = 562; 1235)	9.403^{***} (df = 569; 1228)	229.461^{***} (df = 1; 2550)	13.112^{***} (df = 740; 1811)	13.673^{***} (df = 747; 1804)

Panel B: Carbon Emission is the main explanatory variable

	S&P400			S&P500		
	(1)	(2)	(3)	(1)	(2)	(3)
Emission Level	0.002	0.004	-0.003	0.015^{***}	0.055^{***}	0.047^{***}
	(0.005)	(0.014)	(0.013)	(0.005)	(0.016)	(0.016)
Constant	3.450^{***}	3.026^{***}	3.020^{***}	4.028^{***}	2.601^{***}	3.204^{***}
	(0.123)	(0.340)	(0.337)	(0.118)	(0.505)	(0.514)
Stock FE	NO	YES	YES	NO	YES	YES
Time FE	NO	NO	YES	NO	NO	YES
Observations	1,798	1,798	1,798	2,182	2,182	2,182
R^2	0.0001	0.791	0.801	0.004	0.892	0.897
Adjusted R^2	-0.0005	0.696	0.708	0.004	0.861	0.866
Residual Std. Error	0.563 (df = 1796)	0.311 (df = 1235)	0.304 (df = 1228)	0.677 (df = 2180)	0.253 (df = 1691)	0.248 (df = 1682)
F Statistic	0.130 (df = 1; 1796)	8.310^{***} (df = 562; 1235)	8.664^{***} (df = 569; 1228)	111.976^{***} (df = 1; 2180)	29.834^{***} (df = 490; 1691)	31.602^{***} (df = 499; 1682)

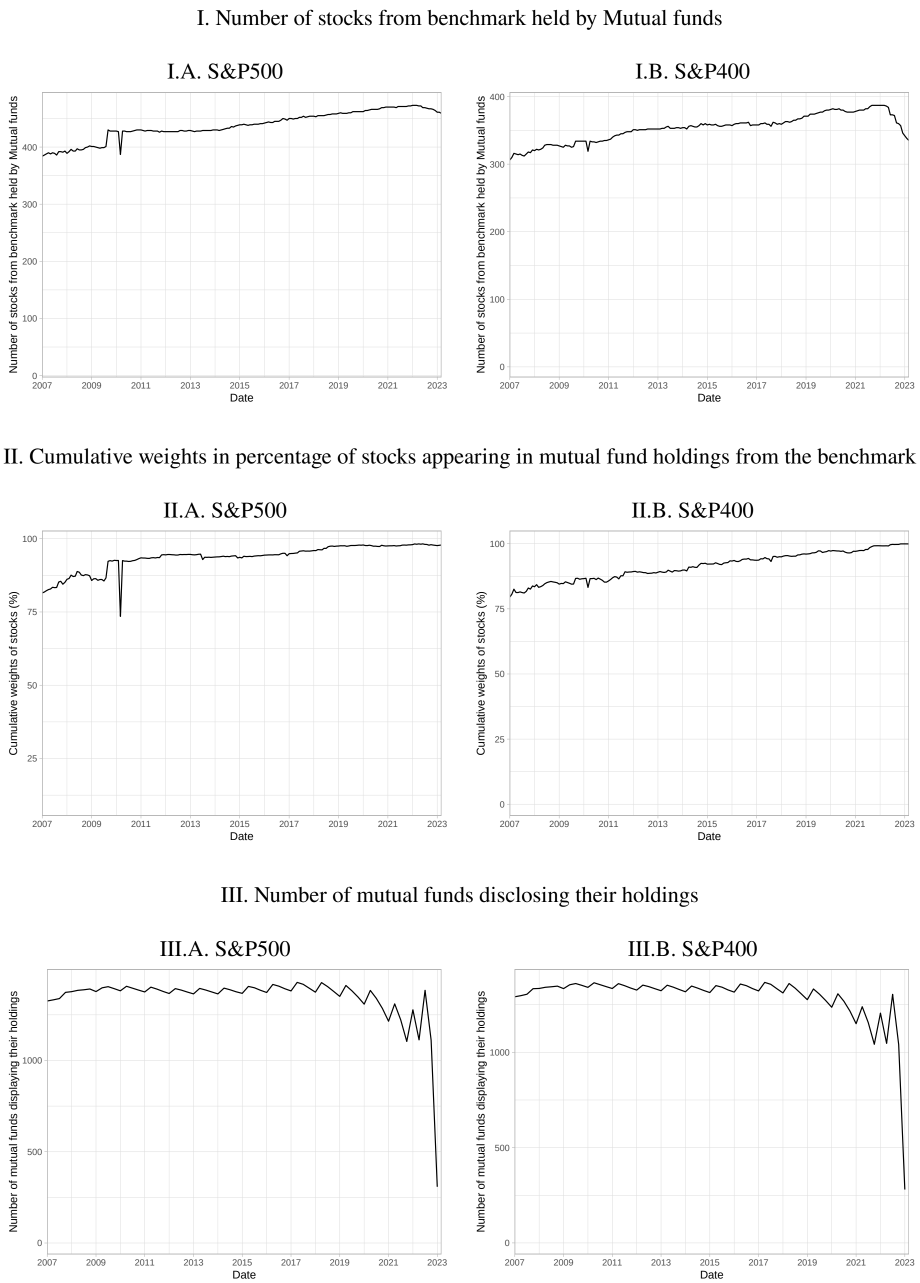

Figure 1: **Stocks' holding and mutual funds portfolios' disclosures through time.**
This figure presents the time series of number of stocks from the benchmark disclosed by at least on mutual fund (I), the cumulative percentage weight of disclosed stocks in the benchmark portfolio (II), and the number of mutual funds disclosing their portfolios with at least one stock from the benchmark (III).

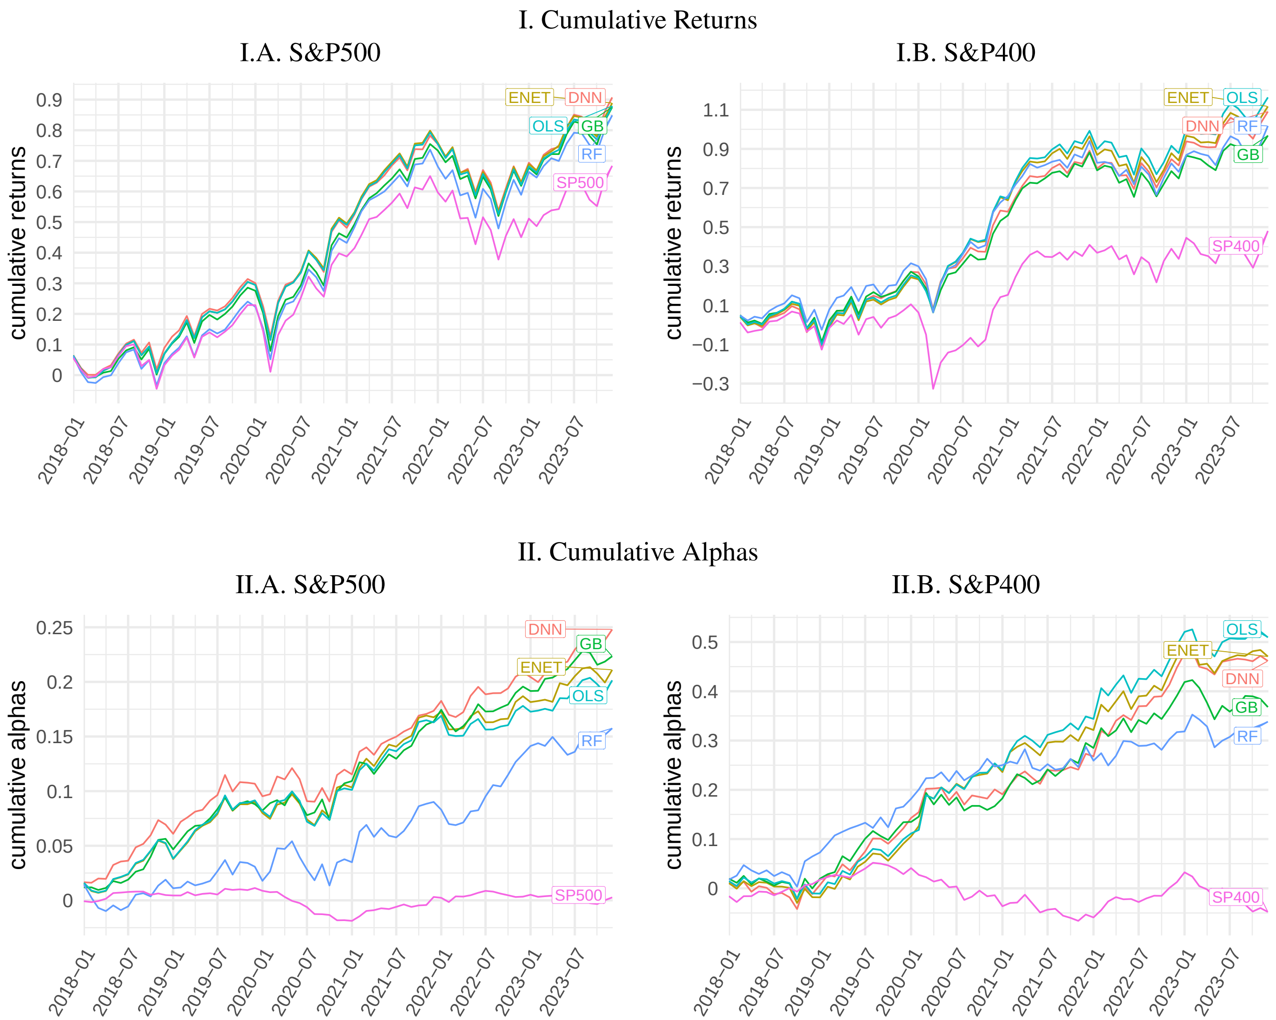

Figure 2: **Out-of-sample (Top portfolios) SAS Strategy Cumulative Returns and Cumulative Alphas - with quarterly rebalancing.** This figure presents the evolution of cumulative returns (I.) and cumulative alphas (II.) of machine learning-based SAS portfolios for different benchmark indices (SP400 and S&P500). We form the machine learning-based SAS portfolios using five different prediction models: OLS, elastic net, random forest, gradient boosting, and deep neural networks. The portfolios are formed every quarter by selecting the top 10 percent of stocks with the highest predicted SAS values. Selected stocks are weighted based on their predicted SAS values in the portfolio. The out-of-sample prediction period goes from January 2018 to December 2023.

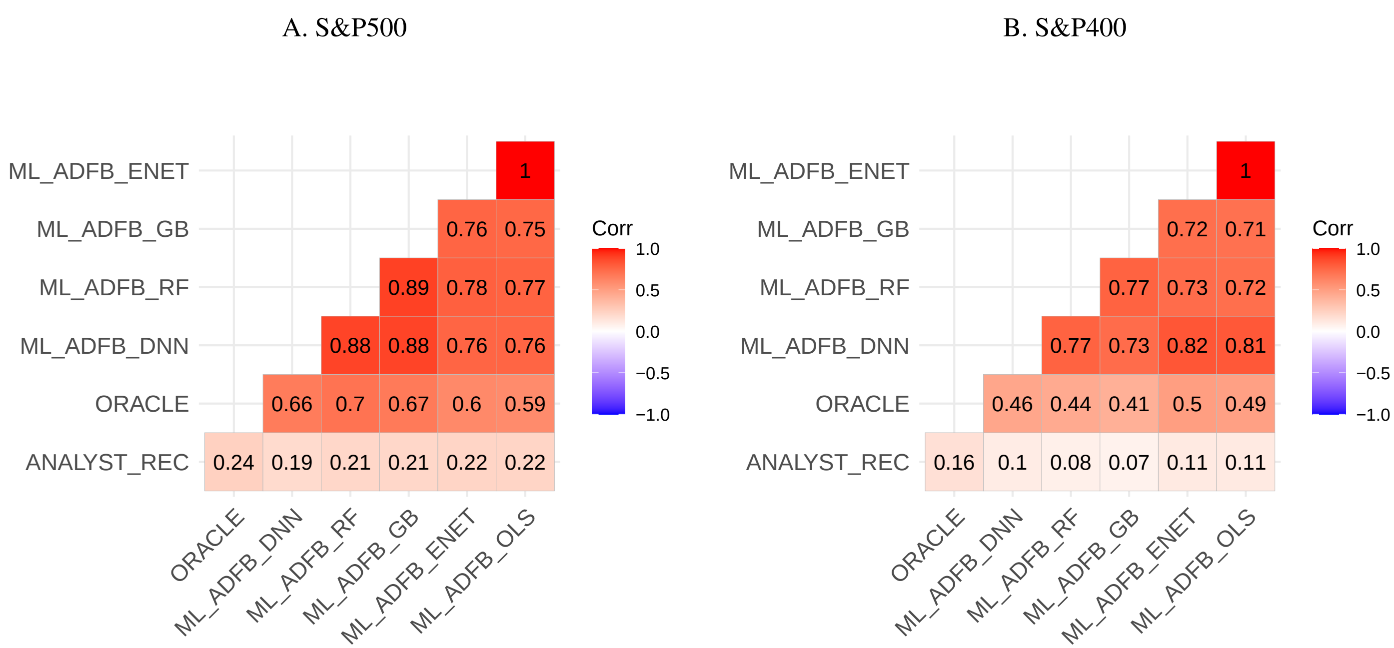

Figure 3: **Correlation Matrices between predicted SAS measures and Analyst Recommendation Consensus - quarterly rebalancing.** This figure presents the correlations between the machine learning-predicted SAS measures, SAS-Oracle measures, and analyst recommendation consensus for two different benchmark indices (SP500 and S&P400). The latter aggregates analyst recommendations for *Strong Buy* (value between $1\&1.49$), *Buy* (value between $1.5\&2.49$), *Hold* (value between $2.5\&3.49$), *Underperform* (value between $3.5\&4.49$), and *Sell* (value between $4.5\&5$). The SAS-Oracle measures (ORACLE) are the real value of stocks SAS values as computed from observed data. The prediction of the SAS measure uses different machine learning algorithms: random forest (RF), gradient boosting (GB), ordinary least squares (OLS), elastic net (ENET), and deep neural network (DNN). Every quarter, we expand the training sample, and the prediction period spans from January 2018 to December 2023. We have tested all the correlations and found that they are statistically different from zero.

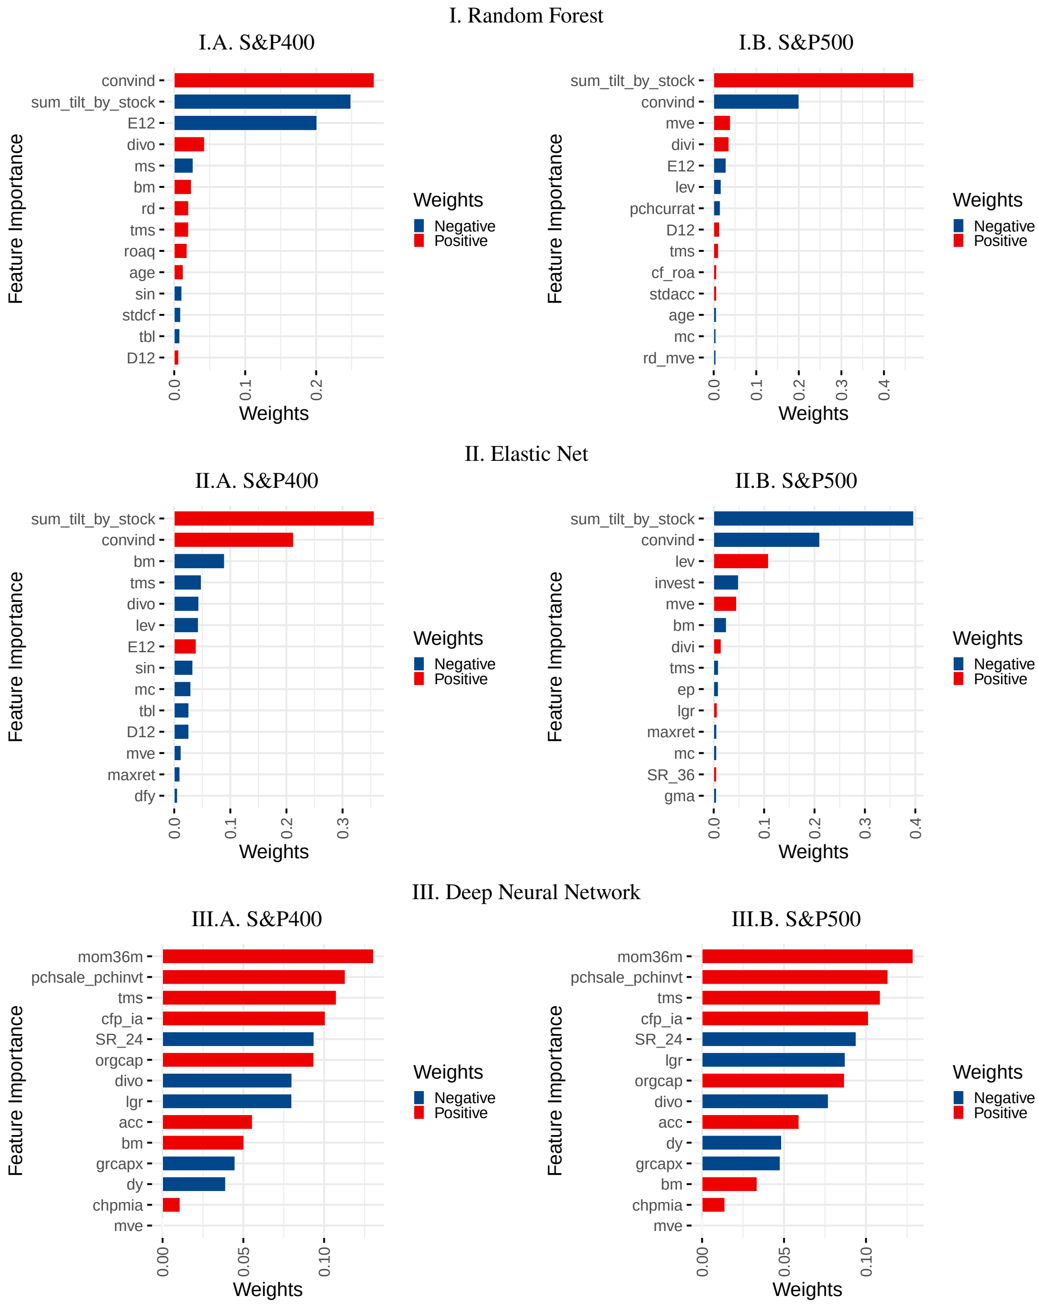

Figure 4: **Feature Importance in the MLSAS strategy with quarterly rebalancing.**
This figure presents aggregate feature importance computed by aggregating the thirty most important features across all entities in the test sample. Roughly the one-third most important variables for making the prediction. Results are displayed for each benchmark index (SP400 and S&P500) and various machine learning models: random forest (I.), elastic net (II.), and deep neural network (III.). Refer to Table 4 for variables definition. *sum_tilt_by_stock* is the lagged dependent variable.

EXTERNAL APPENDIX of **"Learning from the Wisdom of Mutual Fund Managers"**

This supplemental appendix for "Learning from the Wisdom of Mutual Fund Managers" provides additional tables and figures that complement the analysis presented in the main text either with more detailed information about the data used in the main analysis, or by using alternative datasets for robustness checks.

Contents

List of Additional Tables

List of Additional Figures

A Machine Learning Methods

In this section,we present a brief summary of the machine learning methods that we apply to predict stocks' SAS measures and to select stocks to form our investment portfolio. We use supervised machine learning methods, the history of mutual fund holdings, and stock characteristics to learn the decision rule of mutual fund managers. This means a function that mimics manager skills to transform complex information about investor preferences, stock characteristics, and interactions with macroeconomic conditions into a decision about whether or not to include the stock in the mutual fund's portfolio. We aim to build a model that predicts future stocks' SAS values based on current stocks' characteristics. This approach stems from the idea that mutual fund managers used their skills to select these stocks based on publicly available information.

Our supervised learning objective can be addressed both as a regression problem for predicting a stock SAS measure or as a classification problem for predicting a stock selection or not into the SAS strategy's portfolio. The baseline model in the first case can be the ordinary least squares regression (OLS) model of the SAS measure on lagged stock characteristics, while in the second case we could use a logit or probit model to predict the probability of a stock being selected or not as an overperforming stock given its past characteristics. We try OLS as a benchmark model and various machine learning models such as Elastic Net (ENET), Random Forest (RF), Gradient Boosting (GB), and Deep Neural Network (DNN).

Formally, let us consider y the target outcome variable, which in our case can be a continuous variable of stock SAS measure or a binary variable of stock selection into an SAS-Oracle portfolio. We denote by x the vector of predictors (stock characteristics, business cycle or macro variables). We assume that there is an unobservable function f, we want to learn about, used by skilled managers in order to pick stocks based on their characteristics; that means:

$$y = f(x) \tag{A.1}$$

OLS model assumes that f is a linear function of characteristics, and estimates the model's parameters, denoted θ, by minimizing the mean squared errors:

$$\theta^* = \arg\min_{\theta} \quad \sum_{t=1}^{T}\sum_{i=1}^{B_t}\left(y_{i,t} - x_{i,t}^{'}\theta\right)^2 \tag{A.2}$$

where $y_{i,t}$ is the logarithm of stock i measure of absolute deviation from benchmark at time t, $x_{i,t}$ is the vector of predictors, and B_t is the number of stock in the benchmark index at time t. Given the observed values at time t in the test sample of predictors used to train the model, and the estimated parameters in the train sample, the outcome variable is predicted (out-of-sample) as $x_{i,t}^{'}\theta^*$. The predicted values are then sorted to form the investment portfolios.

A.1 Elastic Net

Although the OLS model is simple to build and understand, it has a notable drawback that is resolved by the elastic net model. Undoubtely, the OLS model is susceptible to data overfitting due to its attempt to incorporate every single predictor. This results, especially with large number of predictors, in a model that exhibits strong performance when evaluated with the data it was trained on, but performs poorly when tested on new, unseen data. The elastic net method incorporates regularization parameters to effectively reduce the values of estimated parameters and prioritize predictors that significantly contribute to minimizing mean squared errors. The elatic net model parameters are obtained as follows:

$$\theta^* = \arg\min_{\theta} \quad \sum_{t=1}^{T}\sum_{i=1}^{B_t}\left(y_{i,t} - x_{i,t}^{'}\theta\right)^2 + \lambda\left(\alpha\|\theta\|_1 + (1-\alpha)\|\theta\|_2^2\right) \tag{A.3}$$

Where $\lambda > 0$ and $\alpha \in [0, 1]$ are hyper-parameters set optimally by cross-validation. $\|.\|_1$ and $\|.\|_2$ denote the L_1-norm and L_2-norm respectively. The elastic net objective function ecompasses the least absolute sum of squares operator (LASSO) regression (when $\alpha = 1$) and the ridge regression (when $\alpha = 0$) as special cases.

A.2 Random Forest

OLS and Elastic net both assume a linear function of the predictors to fit the managers conviction about a stock value, but they do not account for possible non-linearities or interactions between predictors. Decision trees model incorporates multiway predictor interactions. It assume that the functional form is as follows:

$$f(x) = \sum_{m=1}^{M} c_m . 1_{(x \in R_m)}, \tag{A.4}$$

where $1_{(.)}$ is the indicator function, and $R_1, ..., R_M$ represent a partition of feature space into M regions based on predictors and split points. M is an hyper-parameter set by cross-validation. The parameter c_m, for $m = 1, ..., M$, is simply estimated as the average of previous realizations y_t such that x_t belongs to R_m as follows:

$$\hat{c}_m = \arg\min_{c_m} \quad \sum_{t=1}^{T} 1\left(x_t \in R_m\right)\left(y_t - c_m\right)^2 = \frac{\sum_{t=1}^{T} 1\left(x_t \in R_m\right) y_t}{\sum_{t=1}^{T} 1\left(x_t \in R_m\right)} \tag{A.5}$$

Decision trees suffer from high variance of the predicted output. Random forest resolved this problem by bootstrapping the original data and averaging predictions across decorrelated decision trees. The decorrelation of decision trees is achieved by using only a random subset of predictors for building the trees each time a split in a tree is considered. This is done for the purpose of generating variability across the boostrap decision trees, thereby making the average of the resulting trees less variable and hence more reliable.

A.3 Gradient Boosting

Similarly to random forests, gradient boosting is based on decision trees. But instead of aggregating the predictions of independent decision trees across bootstrap samples as done by random forest, gradient boosting uses only the original data. It may start with a random guessing prediction function and then sequentially update the prediction function by adding in a shrunken version of the new decision tree to further reduce the pre-existing prediction error. That means we fit a tree using the current residuals rather than the initial outcome variable as the response. As a result, the gradient-boosting method starts with weak decision trees and converges to strong trees by sequentially learning (from the previous tree) how to better fit the data by giving more weight to those observations that are poorly predicted by the current aggregation of trees.

This learning process can result in data overfitting, particularly when the number of sequences is enormous. To address this, cross validation is employed to determine the optimal number of sequence iterations. The iterative enhancement of the prediction function and its sequential update in the boosting decision tree mirror the sequential processing of data in the deep neural network to minimize prediction errors.

A.4 Neural Network Model

A neural network model (NNM) is a combination of layers of nodes where each node linearly combines information in predictor variables into an input supply to an activation function whose outputs are linearly combined to match the observed target variable as closely as possible according to a pre-specified distance.

A deep neural network with L hidden layers is an approximation to $f(x)$ of the form:

$$y \cong g(x;\theta) = \phi^{(L)}\left(\theta_0^{(L)} + \sum_{m=1}^{M_L} \theta_m^{(L)} z_m^{(L)}\right) \tag{A.6}$$

where $\phi^{(L)}(.)$ is an activation function[23] in the output layer, M_L is the number of neurons in the L (last) hidden layer, and $\theta^L \in \mathbb{R}^{M_L+1}$ is the sub-vector of network weights at layer L. For $l \in \{1, ..., L\}$ and $m \in \{1, ..., M_{(l+1)}\}$, $z_m^{(l)}$ is the input information at layer l coming from the

[23] Examples of activation functions include the following: The identity function ($\phi(z) = z$), the sigmoidal function often seen in logit model ($\phi(z) = \frac{1}{1+e^{-z}}$), step function, the rectified linear unit (ReLU, $\phi(z) = \max(0, z)$), the softplus function ($\phi(z) = \log(1 + e^{z})$), etc.

previous layer and defined as follows:

$$z_m^{(l+1)} = \begin{cases} \phi^{(l)}\left(\theta_0^{(l)} + \sum_{n=1}^{M_l} \theta_{n,m}^{(l)} z_n^{(l)}\right) & \text{if } l \in \{2, ..., L\} \\ \phi^{(l)}\left(\theta_0^{(l)} + \sum_{n=1}^{N} \theta_{n,m}^{(l)} x_n\right) & \text{if } l = 1 \end{cases} \tag{A.7}$$

M_l is the width of the network or number of neurons in the layer l. $\theta^l \in \mathbb{R}^{M_l \times (M_{l-1}+1)}$ is the vector of parameters in layer l; in total there are $1 + \prod_{l=1}^{L}(M_l + 1)$ parameters to estimate, and they are chosen in order to minimize a loss (error) function when training the model as follows:

$$\theta^* = \arg\min_{\theta} \quad D\left(y, g(x;\theta)\right) \tag{A.8}$$

Where D denotes a distance, such as the L_2-norm or a divergence metric[24] between the distribution of the observed data and the one predicts by the model. The initialization of the parameter θ for the optimization is made randomly and depends on the random number seed used in the algorithm. Thus, depending on the chosen seed, the model's estimated parameters and prediction may change. Therefore, following Gu et al. (2020), we use an ensemble method in training and evaluating the performance of our neural network. In particular, we use multiple random seeds to initialize neural network estimation and construct predictions by averaging predictions from all networks.

B The MLSAS Investment Strategy with Monthly Rebalancing

Table A2 presents summary statistics on the financial performance of the MLSAS investment strategy but with portfolio rebalancing at every month. The table displays three blocks (Top 10, Top 5, and Top 3) of five columns (RF, GB, OLS, ENET, and DNN) for different machine learning models. Each block corresponds to a quantile (decile, quintile or tercile) portfolio used for the MLSAS investment strategy. For comparison, columns AR and ORA, respectively, provide summary statistics of the performance of similar strategies based on the analysts recommendation consensus (AR) and the observed SAS measures (ORA) over the same investment period. The

[24] Common divergence metrics include Kullback-Leibler divergence and binary cross entropy.

statistics are also provided for the benchmark indices used by a passive investor. Overall, it confirms the patterns previously highlighted, especially regarding the overperformance of the RF and DNN compared to the benchmark indices. However, the performance of the $\mathbf{MLSAS}$ investment strategy is less impressive than in the case of quarterly rebalancing, in particular when compared to the SAS-Oracle strategy. Indeed, we observe a decrease in the monthly average return achieved by the ML-based SAS strategies, an increase in their standard deviations, and a reduction in the Sharpe ratios; whereas the opposite is observed for the SAS-Oracle strategy displaying a higher average return, a lower standard deviation, and a higher Sharpe ratio than with quarterly rebalancing.

The results indicate that Mutual funds often and effectively adjust their portfolio holdings by deviating from the benchmark indices. For an Oracle investor, having real-time access to their portfolio holdings and utilizing that information more frequently (monthly) in the SAS investing strategy allows for superior performance compared to using that information less frequently (quarterly). As a machine learning investor, it is not very profitable to frequently update the portfolio based on predicted information. This is because the new information added to the training sample, when done too often, is insufficient for the machine learning model to correctly update its parameters and accurately predict the future values of stocks SAS measures. Instead, the prediction happens to be more erratic and leads to a lower performance compared to quarterly rebalancing. During quarterly rebalancing, mutual fund managers generate sufficient fresh information that allows the machine learning model to accurately adjust its parameters.

Furthermore, as previously, we observe a similar deterioration in the performance of the strategies as we move from a concentrate (top decile) to a more diversified (top tercile) portfolio holding. In the case of BigCap stocks, the benchmark index overperforms some $\mathbf{MLSAS}$ investment strategies (RF and GB), in terms of Sharpe ratio for the quintile portfolios, and it beats all the ML-based tercile portfolios. In addition, the alpha generated by the $\mathbf{MLSAS}$ strategy is not statistically different than zero, which confirms our view that the formed portfolios' returns are too noisy to be on average positive once we account for the common pricing factors. Given that rebalancing occurs more frequently, implementing this ML-strategy with monthly rebalancing would certainly result in a negative net risk-adjusted return.

Table A3 similarly displays additional performance measures but for the monthly rebalancing strategy. Overall, it corroborates our view that while $\mathbf{MLSAS}$ investment strategy with monthly rebalancing overperforms the benchmark index, the achieved performance remains far below the upper potential displayed by the SAS-Oracle strategy which leaves room for further improvements.

Figure B4 shows the correlations between analyst recommendation consensus and machine learning-based predicted SAS measures for stocks in different benchmark indices over the period from January 2018 to December 2023, when the machine learning training sample is updated every month to predict the $\mathbf{MLSAS}$ value. We see a positive but low correlation between analyst recommendation consensus and the various machine learning-based predicted SAS measures, ranging from 0.1 to 0.3. OLS and Elastic Net predictions appear to be more correlated to analyst recommendations than other non-linear machine learning predictions, and these correlations are highest among BigCap stocks. We also observe that when the training sample updates more frequently (monthly), the correlations between machine learning predicted SAS values and analyst recommendation consensus are either higher or equal to those when the training sample updates quarterly.

Similarly, the AR columns in Table A2 show the performance of the long-only investment strategy based on analyst recommendations, but with portfolio rebalancing every month. The results of the comparative performance between the AR strategy and the benchmark indices are mitigated. For MidCap stocks, the AR strategy performs similarly to the benchmark index with respect to various measures such as average return, standard deviation, and Sharpe ratio. Skewness appears to be slightly worse for the benchmark portfolio than for the AR strategy. On the contrary, the AR strategy for MidCap stocks has a higher kurtosis than the benchmark index. In both cases, risk-adjusted returns are negative and not statistically different from zero. For BigCap stocks, the AR strategy clearly underperforms compared to the benchmark index over the evaluation period. The AR strategy's average return and Sharpe ratio are lower than the benchmark index for all portfolio quantiles, while its standard deviation is larger. Overall, the AR strategy with monthly rebalancing will generates a negative risk adjusted return after transaction cost, and would definitely not be profitable to an investor.

C Additional tables

Table A1: **Performance of the SAS-Oracle strategy - with monthly rebalancing.**
This table shows the performance of the SAS-Oracle strategy implemented by using the real value of stocks SAS values as computed from observed data. The strategy consists of using computed the market-wide absolute deviation from the benchmark for each stock based on its characteristics at time t as if it was known, then sorting the stocks into quantile portfolios based on their computed SAS, and computing the return of the formed portfolios for the next period $t + 1$. *Oracle* portfolios are formed using ex-ante the observed (ex-post) measure of stocks' SAS computed using mutual fund holdings when released to sort stocks into quantile portfolios and to equally weight them. The evaluation period goes from October 2007 to December 2023.

	Decile portfolio										Quintile portfolio					Tercile portfolio			Benchmark
Statistics	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q1	Q2	Q3	Q4	Q5	Q1	Q2	Q3	Index
Panel I: S&P500 benchmark portfolio																			
Mean	-0.309	0.222	0.506	0.837	0.909	1.277	1.250	1.509	1.764	1.778	-0.016	0.676	1.097	1.381	1.772	0.222	1.119	1.644	0.864
Std. Dev.	7.503	6.318	5.914	5.566	5.211	5.225	5.092	4.916	4.651	4.224	6.730	5.642	5.155	4.929	4.354	6.246	5.129	4.525	4.662
Sharpe Ratio	-0.051	0.023	0.073	0.137	0.160	0.230	0.230	0.291	0.363	0.403	-0.013	0.106	0.198	0.265	0.390	0.024	0.203	0.346	0.169
Skew	-0.867	-0.295	0.012	-0.746	-0.495	-0.385	-0.312	-0.107	-0.277	-0.169	-0.569	-0.436	-0.441	-0.231	-0.239	-0.443	-0.446	-0.241	-0.554
Kurtosis	6.428	3.548	3.787	3.094	1.520	1.638	1.472	0.381	0.573	0.280	4.498	3.048	1.531	0.652	0.437	3.998	1.682	0.433	0.810
Adjusted Sharpe Ratio	-0.051	0.023	0.073	0.134	0.158	0.226	0.227	0.289	0.356	0.398	-0.013	0.105	0.195	0.261	0.382	0.024	0.200	0.341	0.166
Alpha	-1.350	-0.665	-0.450	-0.054	-0.018	0.358	0.382	0.600	0.911	0.963	-0.974	-0.246	0.174	0.492	0.940	-0.731	0.222	0.782	-0.018
s.e (alpha)	0.224	0.133	0.155	0.113	0.099	0.101	0.102	0.098	0.104	0.092	0.159	0.113	0.079	0.082	0.081	0.127	0.074	0.060	0.017
Panel II: S&P400 benchmark portfolio																			
Mean	-0.052	0.372	0.387	0.725	0.711	1.110	1.180	1.481	1.951	2.869	0.194	0.561	0.915	1.334	2.432	0.322	0.929	2.059	0.864
Std. Dev.	8.094	7.357	6.542	6.192	6.462	5.981	5.909	5.629	5.539	5.325	7.550	6.237	6.108	5.679	5.319	6.969	5.967	5.386	5.886
Sharpe Ratio	-0.016	0.040	0.048	0.105	0.098	0.173	0.187	0.249	0.338	0.524	0.016	0.078	0.137	0.221	0.443	0.036	0.143	0.368	0.134
Skew	0.282	-0.239	-0.220	-0.357	-0.033	-0.223	-0.444	-0.353	-0.188	-0.021	-0.054	-0.387	-0.178	-0.414	-0.091	-0.128	-0.254	-0.209	-0.567
Kurtosis	6.713	4.213	3.476	1.798	3.240	0.986	1.656	0.774	0.742	0.357	5.231	2.228	1.874	1.167	0.475	4.185	1.816	0.526	2.033
Adjusted Sharpe Ratio	-0.016	0.040	0.048	0.104	0.098	0.172	0.184	0.245	0.334	0.521	0.016	0.077	0.137	0.217	0.438	0.035	0.142	0.362	0.132
Alpha	-1.057	-0.652	-0.571	-0.225	-0.274	0.193	0.202	0.562	1.007	1.924	-0.824	-0.393	-0.035	0.385	1.489	-0.677	-0.022	1.120	-0.067

Table A2: **Investment Performance of Machine Learning Algorithms - with monthly rebalancing.**

This table shows the performance of the machine learning based SAS investment strategy. The Strategy consists of predicting stock future SAS value based on its characteristics at time t, then sorting the stocks into quantile portfolios, and taking a long position in top quantile portfolios composed of stocks with the highest predicted SAS values. The prediction methods are random forest (RF), gradient boosting (GB), ordinary least squares (OLS), elastic net (ENET), deep neural network (DNN). We also use analyst recommendation consensus (AR) to form the portfolio in a similar fashion to the SAS strategy, and the result are presented in the AR columns. ORA columns show the statistics for the SAS-Oracle strategy implemented by using the real value of stocks SAS values as computed from observed data. The formed portfolio is rebalanced in the beginning of each month, and the investment runs from January 2018 to December 2023. The SAS variable used in the machine learning model is the continuous variable of stock market-wide absolute deviation from benchmark that we computed using the mutual funds' holdings of stock.

	Top 10							Top 5							Top 3							Benchmark
Statistics	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	Index
Panel I: S&P500 benchmark portfolio																						
Mean	1.017	1.050	1.189	1.174	1.174	0.975	1.998	1.044	1.022	1.087	1.102	1.084	1.035	1.954	1.008	1.007	1.059	1.045	1.045	0.974	1.816	1.089
Std. Dev.	5.096	5.056	5.039	5.020	4.946	5.993	4.864	5.248	5.137	5.140	5.105	5.144	5.931	4.861	5.345	5.297	5.305	5.313	5.266	5.922	5.019	5.213
Sharpe Ratio	0.170	0.178	0.207	0.204	0.207	0.138	0.380	0.171	0.170	0.182	0.187	0.182	0.149	0.371	0.161	0.162	0.172	0.169	0.170	0.139	0.332	0.180
Skew	-0.392	-0.299	-0.282	-0.284	-0.215	-0.206	-0.124	-0.240	-0.223	-0.205	-0.215	-0.164	-0.391	-0.210	-0.245	-0.228	-0.239	-0.240	-0.254	-0.329	-0.147	-0.370
Kurtosis	0.271	-0.178	-0.252	-0.235	-0.174	1.240	-0.364	0.233	-0.109	-0.028	-0.084	0.043	1.303	-0.210	0.181	0.059	0.077	0.065	0.172	0.992	-0.198	-0.206
Adjusted Sharpe Ratio	0.168	0.177	0.205	0.202	0.206	0.137	0.378	0.169	0.169	0.181	0.186	0.181	0.148	0.367	0.160	0.161	0.170	0.167	0.169	0.138	0.329	0.178
Alpha	-0.008	0.027	0.123	0.110	0.151	-0.112	0.959	0.003	-0.018	0.004	0.023	0.031	-0.058	0.951	-0.032	-0.050	-0.015	-0.035	0.014	-0.130	0.793	0.004
s.e (alpha)	0.108	0.093	0.098	0.096	0.095	0.146	0.088	0.111	0.069	0.072	0.071	0.075	0.124	0.068	0.100	0.087	0.091	0.094	0.084	0.104	0.090	0.028
Panel II: S&P500 benchmark portfolio																						
Mean	1.222	1.245	1.159	1.145	1.161	0.723	2.735	1.126	1.005	1.178	1.163	1.134	0.857	2.361	1.117	1.044	1.176	1.168	1.114	0.825	2.009	0.906
Std. Dev.	6.570	6.574	6.350	6.330	6.462	6.739	5.868	6.476	6.443	6.556	6.528	6.490	6.794	5.871	6.557	6.548	6.597	6.618	6.539	6.941	6.044	6.838
Sharpe Ratio	0.163	0.167	0.159	0.157	0.157	0.085	0.440	0.151	0.133	0.157	0.155	0.152	0.104	0.376	0.147	0.137	0.156	0.154	0.147	0.097	0.307	0.111
Skew	-0.151	-0.122	-0.082	-0.051	-0.089	-0.369	-0.169	-0.097	-0.147	-0.111	-0.126	-0.143	-0.341	-0.049	-0.289	-0.219	-0.143	-0.131	-0.184	-0.352	-0.156	-0.515
Kurtosis	-0.230	-0.081	-0.115	-0.167	-0.221	0.468	-0.258	-0.158	-0.221	-0.196	-0.154	-0.169	0.405	-0.313	0.226	0.027	0.232	0.224	0.209	0.969	-0.210	1.377
Adjusted Sharpe Ratio	0.163	0.166	0.159	0.157	0.156	0.085	0.435	0.150	0.133	0.157	0.155	0.151	0.103	0.376	0.146	0.136	0.155	0.153	0.147	0.097	0.305	0.110
Alpha	0.072	0.145	0.138	0.086	0.151	-0.408	1.667	0.066	-0.072	0.103	0.092	0.070	-0.229	1.294	0.031	-0.066	0.067	0.056	0.013	-0.261	0.961	-0.060
s.e (alpha)	0.233	0.197	0.179	0.205	0.246	0.213	0.215	0.181	0.169	0.171	0.170	0.190	0.197	0.162	0.150	0.150	0.167	0.167	0.160	0.157	0.152	0.131

Table A3: **Additional Investment Performance Measures of Machine Learning Algorithms - with monthly rebalancing.**
This table shows the performance of the machine learning based investment strategy. The Strategy consists of predicting stock future SAS value based on its characteristics at time t, then sorting the stocks into quantile portfolios, and investing in top quantile portfolios composed of stocks with the highest predicted SAS values. The prediction methods are random forest (RF), gradient boosting (GB), ordinary least squares (OLS), elastic net (ENET), deep neural network (DNN). We also use analyst recommendation consensus (AR) to form the portfolio in a similar fashion to the SAS strategy, and the results are presented in the AR columns. ORA columns show the statistics for the SAS-Oracle strategy implemented by using the real value of stocks SAS values as computed from observed data. The formed portfolio is held for one month, it is rebalanced in the beginning of the next month, and the investment runs from January 2018 to December 2023. The SAS variable used in the machine learning model is the continuous variable of stock market-wide absolute deviation from benchmark that we computed using the mutual funds' holdings of stock.

	Top 10							Top 5							Top 3							Benchmark
Statistics	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	RF	GB	OLS	ENET	DNN	AR	ORA	Index
Panel I: S&P500 benchmark portfolio																						
Sortino	0.259	0.275	0.326	0.321	0.331	0.210	0.716	0.266	0.265	0.287	0.295	0.287	0.225	0.682	0.248	0.251	0.267	0.262	0.264	0.209	0.599	0.275
Info.Ratio	-0.009	0.031	0.138	0.129	0.162	-0.077	1.016	0.003	-0.023	0.005	0.029	0.034	-0.054	1.090	-0.037	-0.058	-0.018	-0.039	0.016	-0.137	0.873	0.016
Drawdown	0.252	0.260	0.271	0.271	0.243	0.258	0.154	0.222	0.226	0.245	0.242	0.223	0.289	0.162	0.225	0.223	0.227	0.229	0.223	0.276	0.171	0.243
VaR	0.104	0.098	0.099	0.099	0.096	0.128	0.079	0.104	0.102	0.102	0.100	0.103	0.129	0.087	0.108	0.107	0.106	0.107	0.108	0.130	0.090	0.103
Turnover	0.983	0.959	0.941	0.930	0.953	0.868	0.888	0.931	0.889	0.890	0.885	0.913	0.758	0.833	0.836	0.806	0.796	0.793	0.823	0.634	0.750	1.147
Panel II: S&P400 benchmark portfolio																						
Sortino	0.258	0.267	0.254	0.252	0.250	0.123	0.861	0.238	0.206	0.249	0.245	0.238	0.153	0.725	0.226	0.210	0.245	0.242	0.228	0.142	0.544	0.161
Info.Ratio	0.035	0.084	0.077	0.046	0.075	-0.233	0.728	0.039	-0.043	0.063	0.057	0.044	-0.159	0.753	0.022	-0.046	0.048	0.040	0.010	-0.214	0.657	-0.054
Drawdown	0.220	0.237	0.230	0.220	0.239	0.365	0.153	0.241	0.255	0.237	0.238	0.241	0.350	0.167	0.272	0.280	0.266	0.266	0.263	0.350	0.211	0.354
VaR	0.132	0.133	0.128	0.125	0.135	0.161	0.103	0.131	0.130	0.135	0.134	0.136	0.163	0.095	0.141	0.138	0.134	0.134	0.140	0.163	0.111	0.153
Turnover	1.626	1.532	1.475	1.466	1.523	1.209	1.404	1.410	1.348	1.298	1.299	1.348	1.090	1.255	1.177	1.129	1.087	1.086	1.120	0.918	1.055	1.385

Table A4: **Test of difference in Performance between investment strategies based on SAS and Analysts recommendations - with quarterly rebalancing.**

This table compares the performance of the SAS-based (oracle and machine learning predictions) investment strategies to the one based on analyst recommendations. The prediction methods are random forest (RF), gradient boosting (GB), ordinary least squares (OLS), elastic net (ENET), deep neural network (DNN). We compute and test the difference in mean return, Sharpe ratio, and alpha (using the Fama-French 5 factor model + Momentum) between investment strategies based on the previously cited methods and on analyst recommendation consensus. ORA columns show the statistics for the SAS-Oracle strategy implemented by using the real value of stocks SAS values as computed from observed data. ORA_1, ORA_2, and ORA_3 are respectively for the SAS-Oracle strategies implemented by using the real value of stocks SAS lagged by one, two, and three months respectively. The formed portfolio is held for one quarter, it is rebalanced in the beginning of the next quarter, and the investment runs from January 2018 to December 2023.

	Top 10									Top 5									Top 3									Benchmark
Statistics	RF	GB	OLS	ENET	DNN	ORA	ORA_1	ORA_2	ORA_3	RF	GB	OLS	ENET	DNN	ORA	ORA_1	ORA_2	ORA_3	RF	GB	OLS	ENET	DNN	ORA	ORA_1	ORA_2	ORA_3	Benchmark
Panel I: S&P500 benchmark portfolio																												
Mean (diff)	-0.038	-0.028	-0.028	-0.014	0.007	0.343	-0.051	-0.062	-0.019	0.050	0.071	0.041	0.049	0.076	0.311	-0.005	-0.042	-0.003	0.003	0.071	0.042	0.048	0.034	0.240	-0.068	-0.103	-0.055	-0.236
t-stat	-0.172	-0.112	-0.117	-0.059	0.028	1.357	-0.207	-0.248	-0.083	0.297	0.428	0.246	0.291	0.451	1.625	-0.026	-0.222	-0.016	0.028	0.585	0.334	0.385	0.277	1.543	-0.453	-0.691	-0.385	-1.432
p-value	0.864	0.911	0.907	0.953	0.978	0.179	0.837	0.805	0.934	0.768	0.670	0.806	0.772	0.653	0.109	0.979	0.825	0.987	0.978	0.560	0.740	0.701	0.782	0.128	0.652	0.492	0.702	0.157
Sharpe ratio (diff)	0.025	0.035	0.038	0.041	0.041	0.116	0.039	0.040	0.042	0.039	0.044	0.038	0.038	0.045	0.104	0.037	0.034	0.040	0.020	0.038	0.031	0.032	0.030	0.079	0.016	0.016	0.022	-0.027
t-stat	0.747	0.964	1.030	1.088	1.119	3.004	1.123	1.095	1.162	1.205	1.507	1.268	1.289	1.454	3.373	1.305	1.186	1.452	1.109	1.895	1.679	1.711	1.581	3.928	0.859	0.829	1.163	-0.944
p-value	0.455	0.335	0.303	0.276	0.263	0.003	0.261	0.273	0.245	0.228	0.132	0.205	0.198	0.146	0.001	0.192	0.236	0.146	0.268	0.058	0.093	0.087	0.114	0.0001	0.390	0.407	0.245	0.345
Alpha (diff)	-0.041	0.024	-0.006	0.008	0.045	0.389	-0.008	-0.039	0.008	0.047	0.080	0.042	0.046	0.095	0.377	0.027	-0.003	0.048	0.032	0.122	0.087	0.089	0.084	0.330	0.006	-0.024	0.018	-0.214
t-stat	-0.279	0.147	-0.037	0.049	0.283	2.093	-0.044	-0.210	0.044	0.361	0.722	0.347	0.385	0.725	3.047	0.195	-0.024	0.383	0.347	1.489	0.964	0.976	1.018	3.303	0.062	-0.262	0.195	-1.807
p-value	0.781	0.883	0.971	0.961	0.778	0.040	0.965	0.835	0.965	0.719	0.473	0.730	0.701	0.471	0.003	0.846	0.981	0.703	0.729	0.142	0.339	0.333	0.313	0.002	0.951	0.794	0.846	0.076
Panel I: S&P400 benchmark portfolio																												
Mean (diff)	0.312	0.287	0.562	0.524	0.494	0.695	0.024	0.029	0.003	0.237	0.275	0.347	0.342	0.297	0.423	-0.050	0.011	0.005	0.185	0.197	0.250	0.253	0.231	0.197	-0.094	-0.155	-0.088	-0.501
t-stat	1.260	1.125	2.051	1.918	1.946	2.256	0.079	0.092	0.010	1.311	1.389	1.834	1.776	1.485	1.727	-0.206	0.047	0.024	1.379	1.311	1.906	1.868	1.727	0.948	-0.564	-0.858	-0.492	-2.464
p-value	0.212	0.265	0.044	0.059	0.056	0.027	0.938	0.927	0.992	0.194	0.169	0.071	0.080	0.142	0.089	0.838	0.962	0.981	0.172	0.194	0.061	0.066	0.089	0.347	0.574	0.394	0.624	0.016
Sharpe ratio (diff)	0.055	0.065	0.097	0.091	0.094	0.133	0.023	0.020	0.020	0.044	0.056	0.063	0.063	0.060	0.087	0.026	0.025	0.025	0.041	0.043	0.048	0.048	0.048	0.056	0.009	-0.001	0.012	-0.050
t-stat	1.385	1.666	2.132	1.965	2.287	2.871	0.514	0.422	0.460	1.470	1.911	2.134	2.087	1.972	2.565	0.897	0.764	0.775	1.975	2.210	2.367	2.243	2.488	2.333	0.474	-0.042	0.607	-1.024
p-value	0.166	0.096	0.033	0.049	0.022	0.004	0.607	0.673	0.645	0.142	0.056	0.033	0.037	0.049	0.010	0.370	0.445	0.438	0.048	0.027	0.018	0.025	0.013	0.020	0.636	0.967	0.544	0.306
Alpha (diff)	0.283	0.405	0.574	0.552	0.555	0.758	0.052	0.151	0.038	0.177	0.289	0.318	0.318	0.286	0.427	0.029	0.052	0.055	0.183	0.245	0.247	0.244	0.239	0.267	0.004	-0.049	0.0002	-0.369
t-stat	0.959	1.709	2.280	2.164	2.483	3.473	0.258	0.558	0.169	0.844	1.716	2.166	2.285	1.762	2.628	0.197	0.308	0.323	1.388	2.126	2.280	2.194	2.293	2.354	0.041	-0.411	0.002	-2.446
p-value	0.341	0.092	0.026	0.034	0.016	0.001	0.797	0.579	0.866	0.402	0.091	0.034	0.026	0.083	0.011	0.844	0.759	0.748	0.170	0.038	0.026	0.032	0.025	0.022	0.967	0.682	0.999	0.017

D Additional figures

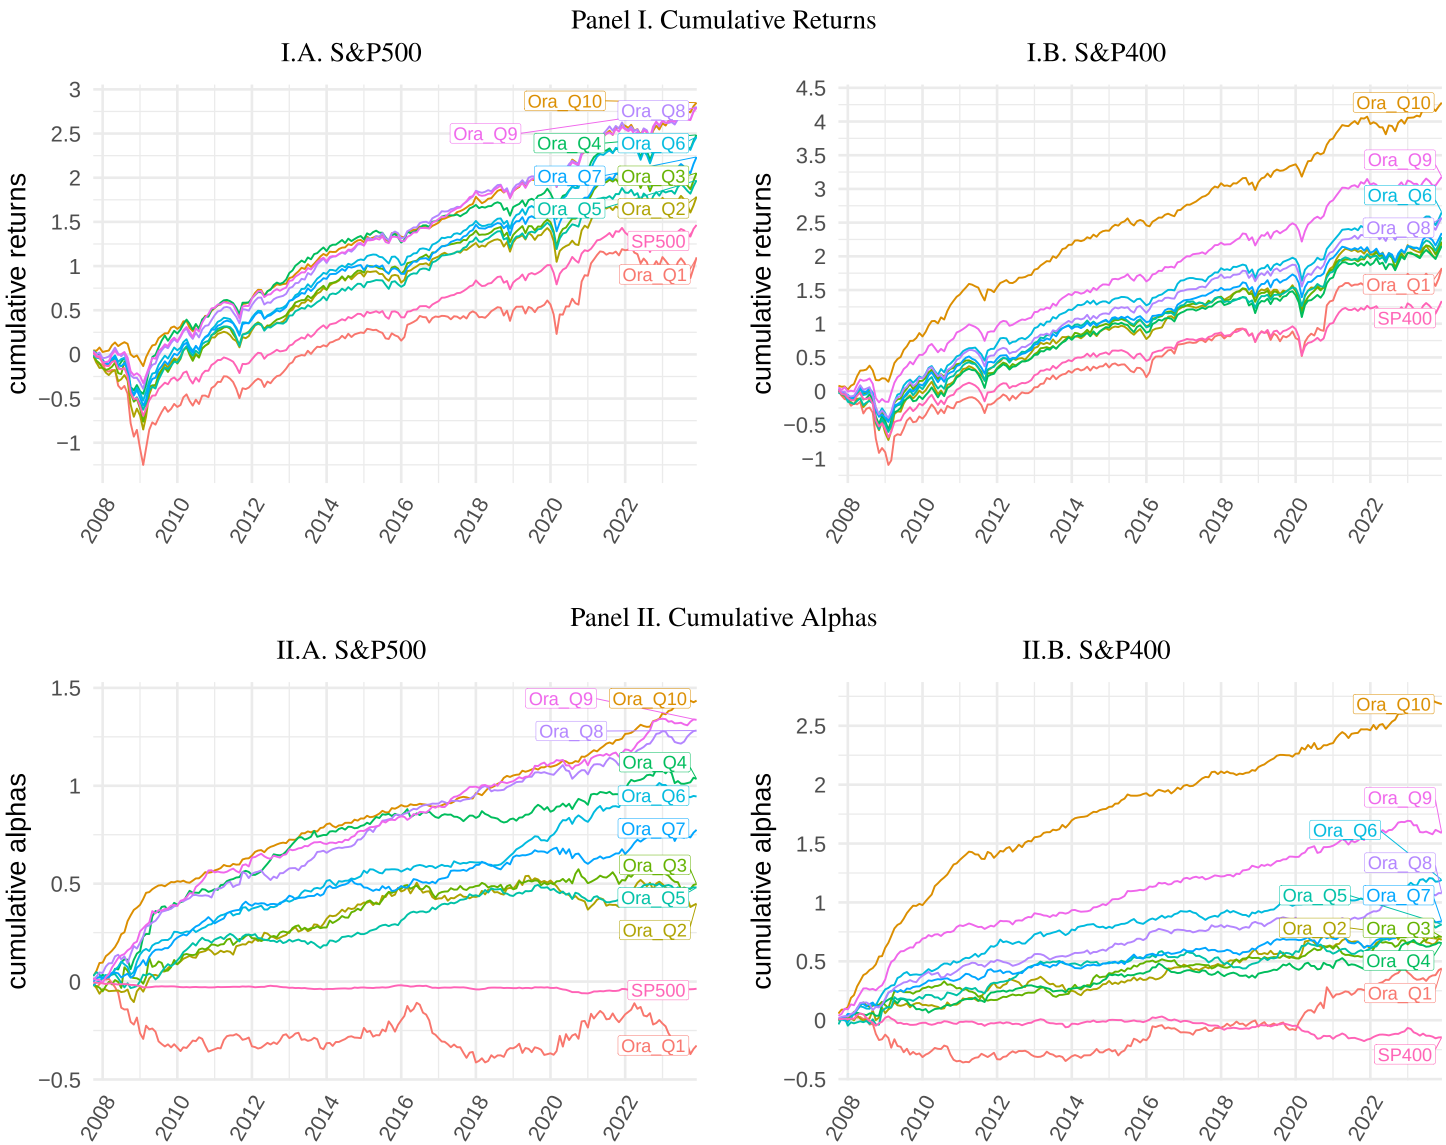

Figure B1: **Historical Cumulative Returns and Cumulative Alphas of the SAS-Oracle Strategy - with quarterly rebalancing.**
This figure presents the evolution of cumulative returns (I.) and cumulative alphas (II.) of SAS-Oracle portfolios for different benchmark indices (SP500 and S&P400). We form the SAS-Oracle portfolios every quarter by sorting stocks into decile portfolios based on their previous month measure of absolute deviation from benchmark (SAS). Selected stocks are weighted based on their SAS values in the portfolio. The evaluation period goes from October 2007 to December 2023.

A. S&P500

B. S&P400

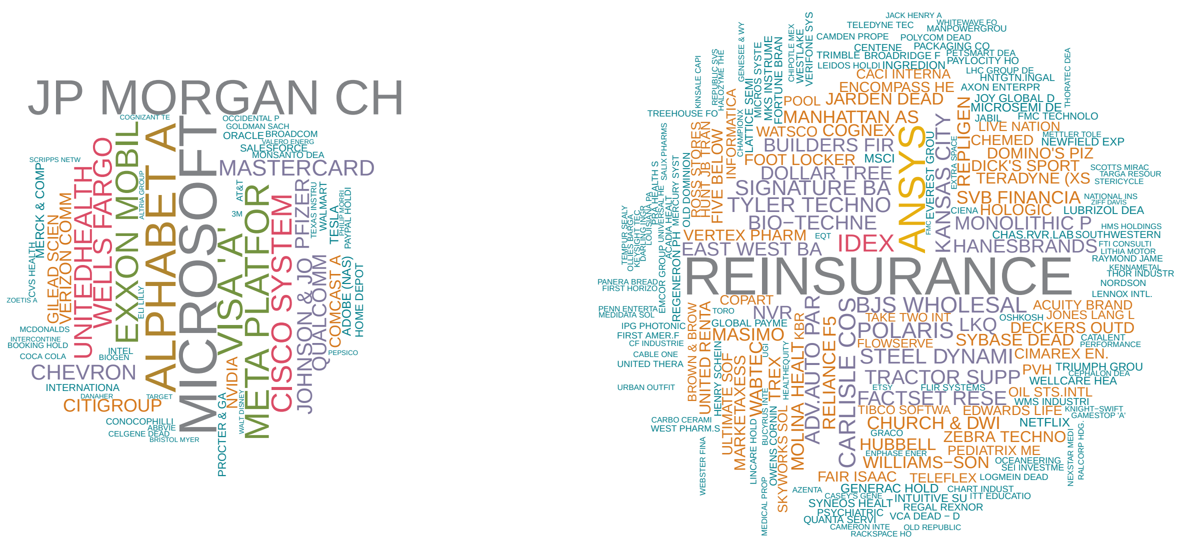

Figure B2: **Stocks in the SAS portfolio by Benchmark index.**
This figure presents a word cloud of the long-only SAS-Oracle strategy that consists of taking a long position in the 5% stocks with the highest SAS values for the S&P500 and S&P400 indices over the investment period. Oracle investors behave as if mutual fund holding information were known in real time. The evaluation period goes from October 2007 to December 2023.

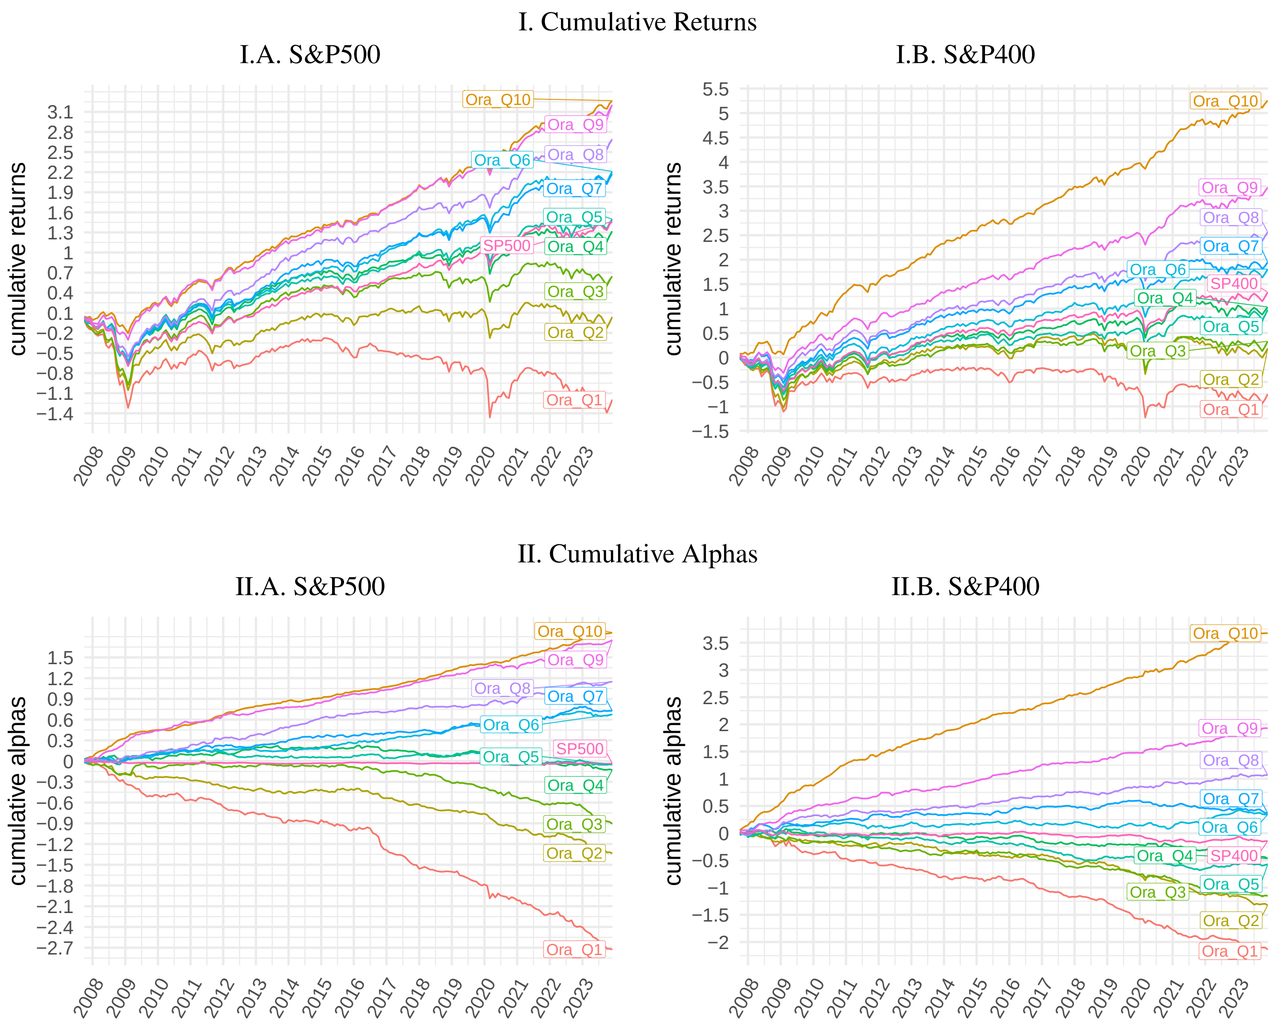

Figure B3: **Historical Cumulative Returns and Cumulative Alphas of the SAS-Oracle Strategy - with monthly rebalancing.**
This figure presents the evolution of cumulative returns (I.) and cumulative alphas (II.) of SAS-Oracle portfolios for different benchmark indices (S&P400 and S&P500). We form the SAS-Oracle portfolios every month by sorting stocks into decile portfolios based on their previous month measure of absolute deviation from benchmark (SAS). Selected stocks are weighted based on their SAS values in the portfolio. The evaluation period goes from October 2007 to December 2023.

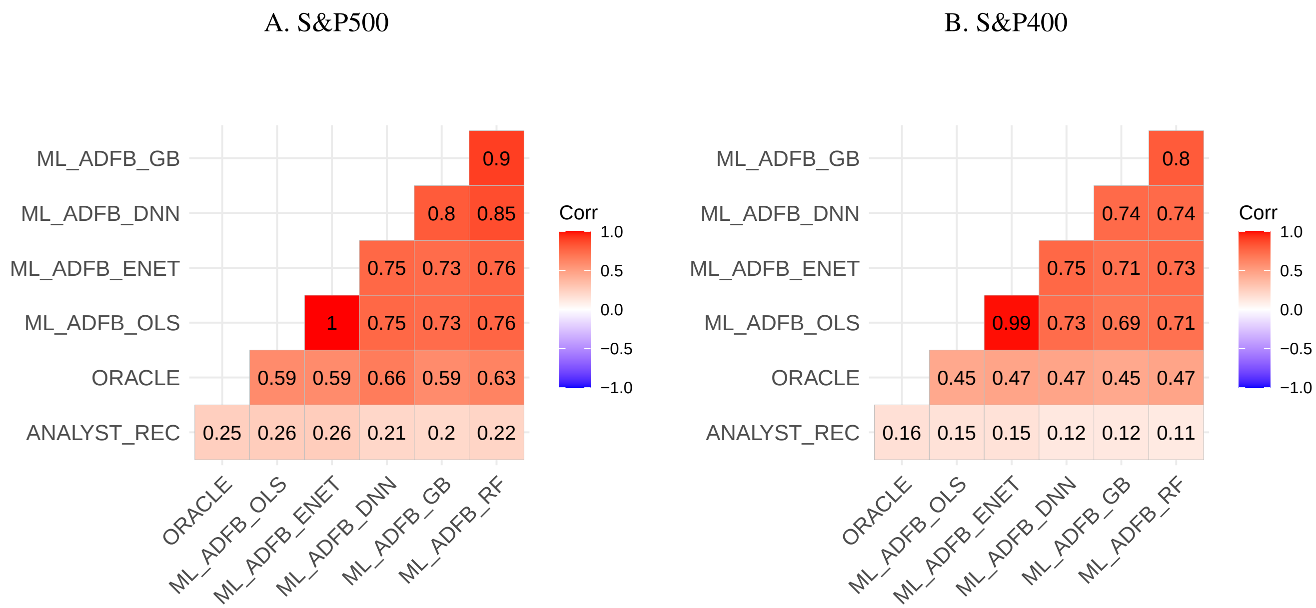

Figure B4: **Correlation Matrices between predicted SAS measures and Analyst Recommendation Consensus - monthly rebalancing.**

This figure presents the correlations between the machine learning-predicted SAS measures, SAS-Oracle measures, and analyst recommendation consensus for two different benchmark indices (S&P500, and S&P400). The latter aggregates analyst recommendations for *Strong Buy* (value between $1\&1.49$), *Buy* (value between $1.5\&2.49$), *Hold* (value between $2.5\&3.49$), *Underperform* (value between $3.5\&4.49$), and *Sell* (value between $4.5\&5$). The SAS-Oracle measures (ORACLE) are the real value of stocks SAS values as computed from observed data. The prediction of the SAS measure uses different machine learning algorithms: random forest (RF), gradient boosting (GB), ordinary least squares (OLS), elastic net (ENET), and deep neural network (DNN). Every month, we expand the training sample, and the prediction period spans from January 2018 to December 2023. We have tested all the correlations and found that they are statistically different from zero.

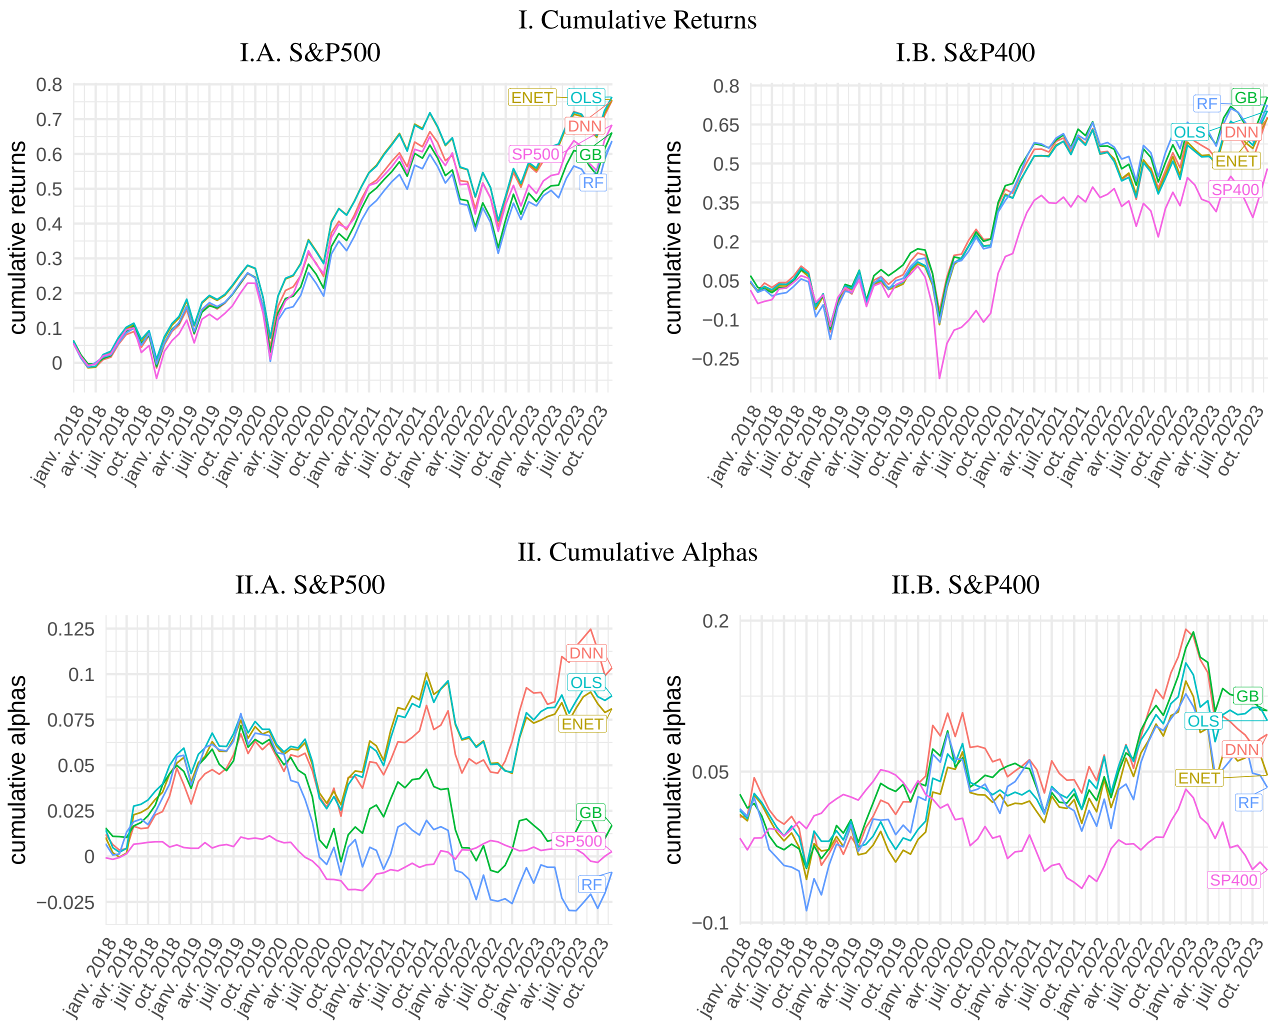

Figure B5: **Out-of-sample (Top portfolios) SAS Strategy Cumulative Returns and Cumulative Alphas - with monthly rebalancing.** This figure presents the evolution of cumulative returns (I.) and cumulative alphas (II.) of machine learning-based SAS portfolios for different benchmark indices (SP400 and S&P500). We form the machine learning-based SAS portfolios using five different prediction models: OLS, elastic net, random forest, gradient boosting, and deep neural networks. The portfolios are formed every month by selecting the top 10 percent of stocks with the highest predicted SAS values. Selected stocks are weighted based on their predicted SAS values in the portfolio. The out-of-sample prediction period goes from January 2018 to December 2023.