Mutual Funds in the Age of AI

September 7, 2024

Abstract

This paper studies the impact of AI technology on the mutual fund industry. I develop a new measure of AI adoption based on hiring practices and find that this measure can predict fund performance. The funds with high AI ratio outperform non-AI funds, after I controlling for standard factors and fund characteristics. Further empirical evidence shows that funds with a high AI ratio tilt their portfolios toward high information intensity stocks, indicating that mutual funds benefit from AI technology adoption by improving their information capacity. Consistent with this channel, I find that the outperformance of these mutual funds mainly comes from better stock picking skills. Finally, AI technology adoption has a negligible effect on fund manager turnover.

1 Introduction

Information is the cornerstone of success in the mutual fund industry, as the fundamental task of a fund manager is to process information about future asset values and use that information to invest in high-value assets. However, fund managers are often overwhelmed by vast amounts of data–from economic indicators and corporate earnings to geopolitical events and market sentiment–yet they have limited time and cognitive capacity to analyze it all effectively (Kacperczyk et al., 2016). This is where artificial intelligence (AI), one of the most significant technological breakthroughs of the past decade, plays a transformative role. Leading asset management firms have heavily invested in AI recently, as evidenced by the rapid increase in human capital investment since 2016 (Figure 2). AI assists fund managers in managing information overload by rapidly processing, analyzing, and extracting insights from enormous datasets. Machine learning algorithms can identify patterns and correlations that might escape human detection, while natural language processing can sift through news articles, social media, and corporate reports to gauge market sentiment and flag potential risks or opportunities.

While past literature has explored the impact of AI adoption on firm growth and innovation (Babina et al., 2024), corporate investment (Sheng-Syan Chen and Peng, 2024), banks (Leonardo Gambacorta and Schiaffi, 2024), hedge funds (Sheng et al., 2024), and venture capitalists (Bonelli, 2023), little is known about its influence on the mutual fund industry. The only indirect evidence comes from Bonelli and Foucault (2023), who find that the release of satellite imagery data can affect the stock-picking skills of mutual fund managers. One potential challenge is the lack of a clear AI measure, as mutual fund companies are not required to disclose this type of information.

This paper provides one of the first pieces of evidence on how AI technology shapes the mutual fund industry. I provide empirical evidence that funds adopting AI technology outperform through portfolio sorting and Fama-MacBeth regressions. Then, I explore the underlying channel, finding that mutual funds with AI technology can gain a competitive

advantage by improving their information capacity. Finally, I also study the impact of AI on mutual fund manager turnover.

To study the adoption of AI technology by mutual funds, I construct a new measure based on hiring practices.[1] The highly specialized nature of AI and its applications demand specialized talent, leading to a scarcity of human capital in this field.[2] This makes the approach particularly appropriate for assessing AI adoption. I collect job posting data for asset management companies from Burning Glass, which encompasses the near-universe of US online job vacancy postings and their detailed skill requirements. Following Babina et al. (2024) and Abis and Veldkamp (2024), I measure the AI-relatedness of each skill in the job postings data by examining its co-occurrence with the four core AI skills. I then calculate the AI-relatedness of each job posting by averaging the AI-relatedness of all the skills required for that position and aggregate this measure to the firm level. Finally, I calculate the AI ratio for each firm-quarter by dividing the AI labor stock by the total labor stock.

Even with the method described above, identifying AI adoption for mutual funds remains challenging due to two potential issues. First, asset management companies typically conduct hiring at the company level, which may include other sectors. For example, Goldman Sachs might hire new employees for its investment banking division. Second, there is a risk of mislabeling, where jobs might be erroneously categorized as AI-related. For instance, asset management firms often hire web developers for website design, and these postings could be mistakenly classified as AI-related due to the programming skills required. To address these challenges, I instructed GPT-4 to act as a judge to determine whether a job posting pertains to the asset management sector and whether it is AI-related. The output from GPT-4 demonstrates that it can efficiently resolve both issues.

I begin my analysis by describing key patterns in AI recruitment for mutual funds. During

[1]The measure is at the firm level rather than the fund level. For more details, see Section 2.4.

[2]See *Wall Street Banks Are Poaching Rival AI Talent* Link: https://www.bloomberg.com/news/articles/2023-11-28/goldman-raided-by-recruiters-in-wall-street-fight-for-ai-talent.
The war for AI talent is heating up Link: https://www.economist.com/business/2024/06/08/the-war-for-ai-talent-is-heating-up

the sample period, the fraction of AI jobs has increased over time, particularly after 2016. Over the entire sample period, the average AI ratio is 1.369%, indicating that for every 1,000 employees in asset management companies, there are, on average, 14 AI employees. The AI ratio is slightly positively correlated with the flow but not correlated with the expense ratio, fund age, turnover ratio, or active share.

Next, I address the question of whether AI adoption can affect fund performance. I conduct two standard analyses in the literature: portfolio sorting and Fama-MacBeth regression. The portfolio sorting results show that mutual funds with a higher AI ratio generate higher returns over the next six months. In my baseline results, I demonstrate that a long-short portfolio, which goes long in the top 20% of funds with the highest AI ratio and short in the bottom 20% of funds with the lowest AI ratio, delivers an annual excess return of 1.56%, statistically significant at the 1% level. Moreover, the results are robust when performance is measured using CAPM alphas or Carhart alphas. These findings are further confirmed in a multivariate analysis that controls for fund characteristics. Using Fama-MacBeth regressions, I show that a 1-standard-deviation increase in the AI ratio is associated with an annualized return that is 61.4 basis points higher. Overall, these results support the conclusion that adopting AI technology provides mutual funds with a competitive advantage over other funds.

Then, I examine return predictability across different sample periods. There are two competing hypotheses. On one hand, since AI became more powerful after 2016, we might expect return predictability to be stronger in the second half of the sample period. On the other hand, AI technology might exhibit a "first mover advantage," where early adopters earn more profit initially, and once all funds adopt AI, none may generate excess returns anymore. In this case, we would expect return predictability to be stronger in the first half. I equally divide the entire sample period into two parts. I repeat the portfolio sorting and Fama-MacBeth regressions for these two subsamples. The results indicate that the predictive power comes from the second half of the sample, supporting the first hypothesis. This suggests

that we have not yet reached an equilibrium where AI technology has stopped evolving and mutual funds no longer benefit from using AI due to overcrowding in AI-based investments.

Having established that the AI ratio can predict the future performance of mutual funds, a natural follow-up question concerns the underlying mechanism. Anecdotal evidence and past literature (Bonelli and Foucault, 2023) suggest that AI technology can enhance mutual funds' information capacity. I provide empirical evidence supporting this specific channel while not excluding other potential channels (such as algorithmic trading) in my analysis. Cao et al. (2021) show that AI surpasses human analysts when the information is transparent but voluminous, whereas humans excel when critical information requires institutional knowledge or subjective judgment. Therefore, my hypothesis is that if mutual funds improve their information capacity after adopting AI technology, they tend to hold high information intensity stocks, in which they have a comparative advantage. I use three variables from Cao et al. (2021) to measure the information intensity of a stock: the number of information events, market capitalization, and age of the stock. Consistent with my hypothesis, I find that mutual funds with a higher AI ratio tend to tilt their portfolios toward larger stocks, older stocks, and stocks with more information events. Furthermore, I find that trades in stocks with high information intensity made by mutual funds with a high AI ratio can predict stock returns in the next quarter, whereas trades in stocks with low information intensity cannot. This evidence supports the conclusion that the outperformance is indeed attributable to high information intensity stocks.

I further provide causal evidence for the above channel using difference-in-difference regressions. I treat the publication of the Transformer model as an AI technology shock and compare mutual funds' holdings before and after this event. I find that mutual funds with a higher AI ratio tilt their portfolios toward stocks with high information intensity following the AI technology shock, which is consistent with the hypothesis.

To better understand the source of outperformance for these mutual funds with a high AI ratio, I decompose the fund performance into stock picking and market timing, following

Kacperczyk et al. (2014)'s method. The results show that the outperformance mainly comes from improved stock picking skills, which is consistent with the idea that mutual funds with a high AI ratio are better at analyzing stocks with high information intensity.

Finally, I test the impact of AI technology on mutual fund managers. On one hand, if mutual funds increasingly rely on AI technology, they may reduce their dependence on individual fund managers. On the other hand, AI technology may not easily threaten mutual fund managers, as this is a high-tech occupation requiring numerous soft skills. Therefore, whether the AI ratio can predict higher fund manager turnover remains an open question. Following Kostovetsky and Warner (2015), I construct two manager turnover variables as the dependent variables. Both OLS and probit regressions show that the AI ratio cannot predict manager turnover, indicating that AI technology has not yet threatened the positions of mutual fund managers.

My paper contributes to the growing literature on the effects of AI on investment. Recent work make progress in examining the impact of AI technologies on investment across various specific settings, such as stock investment (Cao et al., 2021), corporate investment (Sheng-Syan Chen and Peng, 2024), bank lending (Leonardo Gambacorta and Schiaffi, 2024), and VC investment (Bonelli, 2023). To the best of my knowledge, this paper is the first to develop a measure of AI adoption by mutual funds–a previously unexplored class of financial intermediariesand discuss its consequences. More broadly, following the pioneering research by Gu et al. (2020), many researchers use different machine learning tools to develop investment strategies in the stock market to generate excess returns, such as Avramov et al. (2023), Chen et al. (2024), and Li et al. (2022). Recently, some researchers also develop investment strategies leveraging the recent breakthroughs in large language models (Chen et al., 2022; Gabaix et al., 2023; Kim et al., 2024; Lu et al., 2023; Lopez-Lira and Tang, 2023). This paper examines the same topic from a new perspective: whether mutual funds use and benefit from these AI-driven investment strategies.

My paper also highlights that mutual funds can improve their information capacity by

adopting AI technology. The work of Bonelli and Foucault (2023) is closely related to my research, as it explores how the combination of big data and AI skills enables asset managers to gain more precise insights into stock returns and make better investment decisions. Similarly, using mutual fund holding information, Du et al. (2023) find that humans reallocate their information production capacity towards portfolio firms where they have a comparative advantage over machines. However, both papers focus on mutual funds utilizing specific types of AI tools–satellite imagery for Bonelli and Foucault (2023) and automated downloading of SEC filings for Du et al. (2023). In contrast, my paper examines mutual fund AI usage from a broader perspective and provides empirical evidence that it improves the information capacity of mutual funds.

My paper more broadly contributes to the long-standing literature on fund return predictability. Numerous fund characteristics have been used to predict fund returns, such as fund size, fund family size, and turnover. Recent literature has employed machine learning methods to predict fund returns (Li and Rossi, 2020; Kaniel et al., 2023; DeMiguel et al., 2023). I contribute to this body of work by focusing on a new fund characteristic–AI labor recruitment–which can also predict future fund returns. One article similar to mine is that by Abis (2020), which studies how quantitative investment strategies influence mutual fund performance.

Finally, my paper is related to the growing literature on the competition and threat posed to humans by AI technology. Acemoglu et al. (2022) study the effect of exposure to AI technologies on labor demand. Abis and Veldkamp (2024) examine the shift in labor shares within the financial industry driven by new data management and AI jobs. My paper contributes to the existing literature by testing whether AI technology threatens the positions of mutual fund managers.

The paper proceeds as follows. Section 2 describes an AI measure for mutual funds. Section 3 tests whether this AI measure can predict mutual fund performance. Section 4 further investigates the underlying mechanism for return predictability. Section 5 examines

whether adopting AI technology increases fund manager turnover. Section 6 concludes. The Appendix provides details about the variable definitions and the process of constructing the AI measure.

2 Construct AI Measure

AI is a broad and evolving concept. According to the National Institute of Standards and Technology, an AI system is defined as "an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments."[3] In this section, I provide an overview of how AI technology is used in the mutual fund industry. Then, I briefly discuss how previous literature constructs the AI measure. Finally, I introduce the procedure for developing the AI measure in this paper. In addition, I also give an introduction to Burning Glass data, which is key to constructing the AI measure.

2.1 Institutional Background: AI and Mutual Funds

AI has been one of the most significant technological advancements in the past decade. It has been integrated across various industries, including healthcare, retail, transportation, and entertainment. The finance industry is an early adopter of AI and big data technology. Acemoglu et al. (2022) document that the finance sector ranks third in the number of AI job postings, following the information and business services sectors. Within the finance industry, researchers study the impact of AI on stock market (Dou et al., 2024), entrepreneurship (Gofman and Jin, 2024), and sell-side analysts (Grennan and Michaely, 2020). However, there is limited literature focusing on the impact of AI on mutual funds so far.

In practice, mutual funds usually take advantage of AI technology in several ways. First, mutual funds can use AI tools to gather and analyze information to enhance investment

[3]See https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

decision-making. For example, mutual funds can analyze satellite images to find trading signals and generate excess returns. Bonelli and Foucault (2023) find that mutual funds' stock-picking ability in a given stock drops after it becomes "covered" when the satellite image data becomes available. Large language models, recent breakthroughs in generative AI, can also be used to analyze information and make investment decisions. Lu et al. (2023) use ChatGPT to form portfolios based on two types of textual data: Wall Street Journal articles and policy announcements by the Chinese government. They find that the portfolio generated by ChatGPT can significantly outperform the benchmark. Bertomeu et al. (2023) find that after ChatGPT was banned in Italy, the information processing capacity of analysts and investors decreased significantly.[4]

The second application of AI technology in the fund industry is algorithmic trading, often referred to as high-frequency trading (HFT). Leveraging AI's ability to execute trades within milliseconds and handle large volumes simultaneously, algorithmic trading is widely employed to identify small price discrepancies in the market for arbitrage opportunities. In some cases, the processes of information gathering and trading are integrated. Asset managers develop sophisticated mathematical models that analyze market data–such as price, volume, and volatility–to identify trading opportunities. When such opportunities arise, pre-programmed computer algorithms execute trades at high speeds.

Finally, AI is also transforming customer service in asset management. AI-powered chatbots deliver continuous support, adeptly handling queries and issues around the clock. This capability allows for the efficient management of routine inquiries, thereby reducing operational costs. Another application is robo-advisory, which uses algorithms and machine learning to provide automated, low-cost investment advice and portfolio management services to clients. DAcunto et al. (2019) find that investors adopting robo-advising exhibit declines in

[4]The asset management company adopts this kind of textual analysis method to generate trading signals even before ChatGPT was developed. For example, in 2019, BlackRock used technology to analyze over 5,000 earnings call transcripts and more than 6,000 broker reports every day, transforming unstructured text into proprietary measures of trending analyst sentiment. See https://www.blackrock.com/corporate/literature/whitepaper/viewpoint-artificial-intelligence-machine-learning-asset-management-october-2019.pdf.

behavioral biases and experience diversification benefits.

2.2 AI Measures in Literature

Measuring AI usage is difficult since companies are not required to disclose this type of information. There are several different ways to measure AI technology adoption in the previous literature. The first one is based on firms' earnings conference calls. Sheng-Syan Chen and Peng (2024) use textual analysis to capture references to AI applications within management presentations and their responses during Q&A sessions. Similarly, Abis (2020) conduct textual analysis on "Principal Investment Strategies" section of mutual fund prospectuses to categorize funds as quants or discretionaries. However, this type of method cannot be applied to identify AI funds. Chen and Ren (2022) try to identify mutual funds adopting AI technology by analyzing the prospectus (filed as Form 497K or 485BPOS). But they only find 15 AI-powered mutual funds.[5] Researchers also measure AI adoption using survey data. Leonardo Gambacorta and Schiaffi (2024) identify "AI banks" using information obtained from the 2022 RBLS survey. However, such survey data is not available for the mutual fund industry. Another measure is automated information acquisition. Du et al. (2023) use the EDGAR Log File data to infer algorithm usage. If a large volume of EDGAR filings is downloaded beyond human comprehension within a short period of time, it is classified as automation of information acquisition. Then, they identify IP addresses that belong to investment companies. Although machine-based SEC filing downloads are related to AI technology, it is just a simple application and cannot serve as a comprehensive AI measure in my analysis. A recent measure proposed by Sheng et al. (2024) calculates portfolio changes in response to AI-predicted signals from earnings conference call transcripts. This measure also focuses on a specific type of AI usage, whereas I aim to capture the overall use of AI.

The most commonly used measure in recent literature is the intensity of AI-skilled hiring

[5] I also try to identify funds with AI in this way and end up with 19 AI-powered mutual funds by the end of 2022, after excluding funds investing in AI companies. Most of them are active ETFs. The tickers of these mutual funds are: AIVL, AIVI, AQGX, AIEQ, AIIQ, BIKR, QRFT, AMOM, WIZ, HDIV, SNUG, NVQ, DUDE, BOB, LETB, OAIE, AIDB, LQAI, AIYY.

(Acemoglu et al., 2022; Babina et al., 2024; Abis and Veldkamp, 2024; Cao et al., 2022). Given that AI is highly technical and its applications require specialized talent, this approach is particularly suitable. The basic idea is to leverage job postings data to calculate an AI score for each skill and aggregate to the job level and then company level. To construct a recruitment-based AI measure, the key input is Burning Glass job posting data, which will be introduced in the next subsection.

2.3 Burning Glass Data

Job posting data is sourced from Lightcast (formerly Burning Glass Technologies, referred to as Burning Glass hereafter), a premier labor market analytics firm in the United States. Burning Glass aggregates data from a comprehensive range of online sources, including approximately 40,000 company websites and job boards, with no more than 5% of vacancies from any one source. The firm employs a deduplication algorithm to refine the data, transforming it into a format suitable for analysis. Burning Glass data capture the near-universe of jobs posted online and cover 60%-80% of all U.S. job vacancies. The finance and technology industries have particularly good coverage. Besides being useful for job seekers, this data is also widely used by researchers in the field of labor economics. Acemoglu et al. (2022) show that the data closely track the evolution of overall vacancies in the US economy as recorded by the nationally representative Bureau of Labor Statistics (BLS) Job Openings and Labor Turnover Survey (JOLTS), which verifies the representativeness of the Burning Glass data.

My sample includes data spanning from the beginning of 2010 until December 2022, as the Burning Glass data starts from 2010. After removing duplicated job postings, I match the employers listed in the postings with asset management companies. I conduct fuzzy matching between company names in the Burning Glass database and the names of asset management companies in the CRSP mutual fund database. For observations that do not exactly match, I manually assess the top three potential fuzzy matches by examining the company names. I exclude asset management companies that have fewer than 100 job postings due to the

potential for significant noise.[6] The final sample of job postings contains a total of 5,329,188 observations.

2.4 Measuring AI for Mutual Funds

In this subsection, I describe my methodology for measuring AI usage in mutual funds. This methodology is based on those used in Babina et al. (2024), Abis and Veldkamp (2024) and Cao et al. (2022), but includes a few improvements. The steps are as follows: first, I calculate the AI-relatedness of each skill and aggregate it to the job-posting level; second, I adjust the job-posting level AI score using GPT and then aggregate it to the company level; third, I adjust the number of AI job postings based on the estimated hiring/separation rate and calculate the AI labor stock.

The first step is to measure the AI-relatedness of each skill. I basically follow Babina et al. (2024) in this step. Four skills are defined as unambiguous core AI skills: Artificial Intelligence (AI), machine learning (ML), natural language processing (NLP) and computer vision (CV). For each skill s, I calculate their co-occurrence with the core AI skills:

$$w_s^{AI} = \frac{\text{\# of jobs with skill s and (AI, ML, NLP or CV) in required skills or in job title}}{\text{\# of jobs with requiring skill s}} \tag{1}$$

This measure reflects the degree of correlation between each skill s and the core AI skills. I present 20 skills that demonstrate high AI-relatedness and 20 that exhibit low AI-relatedness in Appendix B, Table 12. For instance, the skill "Unstructured Data" has a value of 0.46, indicating that 46% of job postings requiring "Unstructured Data" also require one of the core AI skills or mention one of the core AI skills in the job title. Conversely, "Regulatory Compliance" has a value of only 0.018. These results are consistent with the common sense that "Unstructured Data" is closely related to AI, while "Regulatory Compliance" is

[6] I match roughly half of the mutual fund universe to the Burning Glass database. Most of them are from relatively large fund families.

unrelated.

The next step is to aggregate the AI-relatedness to the job-posting level. In Babina et al. (2024), this is achieved by calculating the average AI-relatedness across all required skills for each job posting. While this approach is generally suitable for most companies, it encounters two potential challenges when applied to asset management companies. First, asset management companies typically conduct hiring at the company level, and the hiring might be for other sectors. For example, Goldman Sachs might hire new employees for its investment banking division rather than its asset management sector. Second, there is a risk of mislabeling, where jobs might be erroneously categorized as AI-related. For example, asset management firms often hire web developers for website design, and these postings could be mistakenly classified as AI-related due to the programming skills required.

To address the two challenges above, I instruct GPT-4 to act as a judge to determine whether a job posting pertains to the asset management sector and whether it is AI-related. Figure 1 illustrates the entire process used to identify AI-related jobs. To give a better interpretation of the procedures, I also show the details for ten examples. Appendix B, Table 13 lists ten job postings from Burning Glass, detailing the company name, required skills, and job title. The AI score is calculated as the average AI-relatedness of all the required skills for job posting j, as in Babina et al. (2024):

$$w_j^{AI} = \frac{1}{N} \sum_{s=1}^{N} w_s^{AI} \tag{2}$$

The first step involves determining whether these job postings are from the asset management sector. I input the job titles into GPT-4 for evaluation. Appendix B, Figure 5 presents the prompt I used and GPT-4's response for the ten job postings listed in Appendix B, Table 13.[7] GPT-4 identifies that the first, second, and sixth job postings are not from the asset management sector, which aligns with our intuition that they are from the banking sector.

[7]This demonstration case is generated by GPT-4 in POE. For the formal empirical analysis, I use the OpenAI API with the same prompt. This setup will not lead to variance in GPT-4's responses because I ask GPT-4 to forget the previous input each time.

After this step, the total number of job postings is reduced from 5,329,188 to 1,853,763. Second, I exclude the job postings with $w_j^{AI} < 0.07$. Here, I try to choose a threshold slightly lower than Babina et al. (2024) (which is 0.1) because the final filtering step by GPT-4 can reduce Type I error (incorrectly labeling other types of job postings as AI-related).[8] After this step, the number of job postings decreases from 1,853,763 to 70,134. The last step involves GPT-4 assessing whether a job posting is AI-related based on its title. Appendix B, Figure 6 presents the prompt I used and GPT-4's response for the seven job postings from the asset management sector in Appendix B, Figure 5. GPT-4 categorizes a senior data scientist as "Strongly AI related", an ESG data specialist as "Weakly AI related", and a lead site reliability engineer as "Not AI related".[9] These results are consistent with iutuition. From the 70,134 job postings remaining after the previous step, 24,123 are classified as "Strongly AI related", 5,049 as "AI related", 22,799 as "Weakly AI related", and 18,163 as "Not AI related". Finally, I categorize the labels generated by GPT-4 into a numerical score using the following mapping:

- $I_j = 1$ is assigned to "Strongly AI related"
- $I_j = 0.7$ is assigned to "AI related"
- $I_j = 0.3$ is assigned to "Weakly AI related"
- $I_j = 0$ is assigned to "Not AI related"

The correlation between this indicator I_j and the raw AI score w_j^{AI} is 67.13%. This correlation suggests that GPT-4's judgments are closely aligned with the AI relatedness of the skills required, meanwhile providing a refined assessment.

Figure 7 illustrates the frequency of all keywords in the titles of job postings categorized as "Strongly AI related" and "Not AI related" by GPT-4, with larger sizes indicating higher

[8] In Appendix C, I also explore other thresholds. I construct alternative AI measures with cutoffs equal to 0.075, 0.08, 0.085, and 0.09. Appendix C, Table 15 and Figure 9 show that the correlations between these measures are higher than 0.99.

[9] In my sample, the AI job postings can be roughly divided into two categories. Some focus on applying AI (e.g., Vice President, Systematic Active Equity Team), while others support AI infrastructure (e.g., Data Scientist - Machine Learning/AI/Python). There are also some in between (e.g., ML Engineer - Investment (Python/AWS)).

frequencies. The keyword with the highest frequency in the "Strongly AI related" category is "big data," while "Java developer" has a high frequency in the "Not AI related" category. This indicates that GPT-4 effectively helps filter jobs with similar skill requirements to AI jobs but are not AI jobs.

After obtaining a measure for AI-related job postings, I aggregate this data to the quarterly level to observe trends in AI hiring over time. Figure 2 plots the AI labor recruitment in the mutual fund industry quarterly. In the upper panel of the figure, AI job postings are relatively scarce during the early years and show a significant increase later on, with the exception of 2020 due to the COVID-19 pandemic. Meanwhile, the total number of job postings in the mutual fund industry exhibits a gradual increase throughout the sample period. The lower panel of Figure 2 plots the ratio of AI job postings to total job postings. This graph highlights a marked surge in AI hiring after 2016Q4. Another takeaway from Figure 2 is that is that despite concerns about AI potentially takes away jobs from people, the mutual fund industry has not yet reached that stage. In Figure 8 in Appendix B, I also show the AI labor recruitment for Blackrock and T. Rowe Price Group as two examples. Their patterns align with the general trend in Figure 2, though the hiring has a larger variation in the company level.

The final step is to calculate the AI labor stock for each asset management company with the indicator above. I follow a similar method to Abis and Veldkamp (2024) and Cao et al. (2022). First, I obtain data from the Bureau of Labor Statistics to estimate the likelihood that a vacancy is filled and the likelihood that an employed worker leaves their job. I compute the labor stock for each firm-quarter as follows:

$$l_{i,t}^{AI} = l_{i,t-1}^{AI}(1 - sep_t^{AI}) + h_t^{AI} \sum_{j=1}^{N} I_{i,j} \tag{3}$$

where $l_{i,t}^{AI}$ denotes the AI labor stock for firm i in quarter t, sep_t^{AI} is the separation rate, h_t^{AI} represents the vacancy fill rate for the financial services sector[10] and $I_{i,j}$ is the indicator

[10]Data is sourced from the Finance and Insurance (NAICS 52) industry according to the BLS classification.

for job posting j at firm i, calculated in the last step. For example, if Firm A has 50 AI employees in 2016Q4 and posts 20 AI job postings in 2017Q1 with an estimated average separation rate of 0.08 and a hiring rate of 0.6 for that quarter, then Firm A's AI labor stock in 2017Q1 would be calculated as follows: $50 \times (1 - 0.08) + 20 \times 0.6 = 58$. Subsequently, I calculate the total labor stock $l_{i,t}^{Total}$ in the same way, using the total number of job postings in the asset management sector for each company. I measure AI usage within each firm by calculating the AI ratio, defined as the ratio of AI labor stock to the total labor.[11]

$$AI_ratio_{i,t} = \frac{l_{i,t}^{AI}}{l_{i,t}^{Total}} \tag{4}$$

It is worth noting that since the hiring is conducted at the firm level, the AI measure is also at the firm level rather than the fund level. The assumption here is that if an asset management company employs a higher proportion of AI labor, the mutual funds it manages tend to utilize more AI on average. In reality, asset management companies often form centralized AI or data science teams. These teams are responsible for developing and maintaining AI models, data analytics, and other technological tools that can be used across the entire organization.

3 AI Adoption and Mutual Fund Performance

In the last section, I documented the rapid adoption of AI technology in the mutual fund industry over the past decade. However, since stock investment is a challenging task, it remains an open question whether mutual funds can benefit from the AI technology. In this section, I test whether the AI ratio can predict mutual fund performance. First, I introduce the data used in the empirical analysis and present descriptive statistics. Then, I conduct portfolio sorting and Fama-MacBeth regressions using the AI ratio constructed in the previous section.

[11] Here, it is important to measure AI hiring as a ratio. Some people may argue that, keeping other conditions unchanged, it is not surprising that increasing the recruitment of a specific type of labor can improve a fund's performance. When I measure AI hiring as a ratio, the question becomes whether hiring more AI employees is a relatively better allocation of human capital, which is an open question.

3.1 Data and Summary Statistics

I use data from a variety of publicly available databases. The first one is the job posting data from Burning Glass, as I discuss in Section 2.3. The second one is the CRSP Survivorship Bias-Free Mutual Fund Database, which contains monthly net returns, total net assets (TNA), and other characteristics (expense ratio, portfolio turnover, fund type, etc.). Net return is the simple return received by the investors after fund expenses. Using the CRSP share class group number (`crsp_cl_grp`), I aggregate the fund return across share classes, value-weighted by TNA. My analysis focuses on actively managed domestic equity mutual funds.[12] I exclude target-date funds by removing funds whose names contain the strings target and specific years (e.g., 2005, 2010, 2015, etc.). I also exclude funds with total assets below $10 million. The third database is the Thomson Reuters Mutual Fund Holdings database (TFN/CDA S12), from which I get quarterly mutual fund holdings information. I use MFLINKS to merge the CRSP Mutual Fund Database and the S12 holdings database. The stock-level information is obtained from the CRSP database, except for the number of information events, which comes from the Capital IQ Key Development database. Market return, risk free rate and Fama-French Charhart four-factor are obtained from Professor Kenneth French's website. The last one is the Morningstar database, from which I get information related to mutual fund managers. The Morningstar database and the CRSP Mutual Fund Database are merged using CUSIP.

Appendix A shows a comprehensive definitions of all the variables in this paper. Fund age is based on the oldest share class. Activeshare is calculated with the method of Doshi et al. (2015).[13] Following the mutual fund literature (e.g. Lou, 2012), the flow rate for fund i in quarter t is defined as the net flow into the fund divided by lagged TNA, adjusted by

[12] A fund is a domestic active equity fund if its CRSP fund style code starts with "ED" and its `index_fund_flag` does not equal "B", "D", or "E".

[13] The code can be found at Professor Mikhail Simutin's website. See http://www-2.rotman.utoronto.ca/simutin/research.asp.

M&A:

$$flow_{i,t} = \frac{TNA_{i,t} - TNA_{i,t-1} * (1 + RET_{i,t}) - MGN_{i,t}}{TNA_{i,t-1}} \tag{5}$$

Finally, I exclude the funds with TNA less than 10 million dollar. I also exclude the data from 2010, as it serves as the formation period for the AI ratio, given that labor stock is calculated cumulatively. All continuous variables are winsorized at the 1% and 99% levels to minimize the impact of outliers.

Panel A of Table 1 presents the summary statistics for my sample. The average AI ratio is 1.369%, indicating that for every 1,000 employees in asset management companies, there are, on average, 14 AI employees. The results for other variables are consistent with those reported in earlier studies, except for TNA and family size, which are relatively larger due to the matching method. Panel B of Table 1 displays the correlation matrix among the main variables. The AI ratio has a low correlation with other variables, while it exhibits a slightly positive correlation with TNA and a slightly negative correlation with the expense ratio. This suggests that funds adopting AI technology are relatively larger and charge lower fees. However, this correlation only serves as weak evidence since I do not add any restrictions or control variables.

After presenting the summary statistics, I document some basic facts about AI funds. I examine the relationship between the AI measure and fund characteristics using the following regressions:

$$Characteristics_{i,t} = \alpha + \beta\, AI_ratio_{i,t-1} + \eta_t + \epsilon_{i,t} \tag{6}$$

where $Characteristics_{i,t}$ represents the fund characteristic of interest for fund i in quarter t. I choose five characteristics: flow rate, expense ratio, fund age, turnover ratio, and active share. The independent variable is the AI ratio, lagged by one quarter. I include time-fixed effects to control for the temporal trend of the AI ratio. Standard errors are clustered at the

fund family and quarter level. Table 16 in Appendix D reports the regression results. Funds with a higher AI ratio tend to have a higher flow rate. Although funds do not disclose AI usage in their prospectus, the integration of AI technologies into investment strategies can be a significant selling point during the marketing process, which can attract inflows. The coefficients for the other four dependent variables are not significant.

Furthermore, I present the distribution of AI ratios across all mutual funds at different points in time to observe how it evolves. The results are shown in Figure 3. The distribution shifts to the right over time, indicating that an increasing number of asset management companies are adopting AI technology.[14]

3.2 Portfolio Sorting

I begin investigating the relationship between the AI ratio and future fund performance using a portfolio sorting method. At the beginning of every semi-year, I sort all mutual funds based on their AI ratio and form quintile portfolios. The high (low) quintile portfolio consists of mutual funds with the highest (lowest) AI ratio values. I conduct portfolio sorting semi-yearly rather than quarterly because it takes time for AI labor to become effective. I then construct a long-short portfolio that goes long on the high quintile portfolio and short on the low quintile portfolio, holding it for one month. Finally, following previous studies, I compute risk-adjusted performance using the CAPM and the Carhart 4-factor model:

$$Alpha_{i,t}^{CAPM} = Ret_{i,t} - \beta_{i,t-1} \times RMRF_t \tag{7}$$

$$Alpha_{i,t}^{Cahart} = Ret_{i,t} - \beta_{i,t-1}^{1} \times RMRF_t - \beta_{i,t-1}^{2} \times SMB_t - \beta_{i,t-1}^{3} \times HML_t - \beta_{i,t-1}^{4} \times MOM_t \tag{8}$$

[14]The mutual funds on the right side of the graphs are managed by AQR Capital Management. AQR (Applied Quantitative Research) is a renowned quantitative investment firm, well-known for its use of quantitative methods to guide its investment strategies. Additionally, AQR is a leading adopter of AI technology in the investment industry.

where $Ret_{i,t}$ is the excess return of fund i in period t over the risk-free rate. $RMRF_t$ is the market excess return, and SMB_t, HML_t, and MOM_t are the returns of the factor portfolios related to size, book-to-market, and momentum, respectively. All βs are calculated using a rolling window regression from $t-36$ to $t-1$. In other words, alpha is defined as the difference between a fund's raw return in period t and the fund's 4-factor expected return in period t.

The first column of Table 2 reports the value-weighted time-series average monthly mutual fund return (in percentage) for funds within each quintile. The next two columns report the value-weighted time-series average CAPM monthly alphas and Carhart 4-factor monthly alphas, respectively. The total number of observations is equal to 144, as the sample period spans from 2011 to 2022, containing 144 months. At the bottom of Table 2, I also report the performance differences in return (alpha) between the portfolios of high-AI (bottom quintile) and low-AI (top quintile) funds. Alphas are negative in most quintiles, which is consistent with the well-documented fact that the mutual fund industry cannot beat the market (Fama and French, 2010). I find that funds hiring more AI employees significantly perform better in the future. Specifically, the difference between the bottom and top quintiles is positive: the monthly performance difference is 0.130% for raw returns, 0.164% for CAPM alphas, and 0.076% for Carhart 4-factor alphas. These translate to annualized return differences of 156 basis points (i.e., 0.130×12) for raw returns, 197 basis points for CAPM alphas, and 91 basis points for Carhart 4-factor alphas. All three performance measures generate differences that are statistically significant. All the results above indicate that mutual funds can benefit from adopting AI technology by outperforming other funds.

After documenting the alphas of the AI ratio quintile portfolios, I also study the differences in factor loadings of these portfolios. The factor loadings are calculated using the daily returns for each fund each quarter. Then, I calculate the average factor loading for each portfolio each quarter. Panel A in Table 6 reports the average factor loading for each group. The factor loadings are similar across AI quintiles, except in the case of the size factor, which is

smaller for portfolios with a higher AI ratio.

3.3 Multivariate Analysis

Next, I perform a multivariate analysis, which allows me to control for a set of fund-specific characteristics that may subsume the AI measure's power to predict fund returns. These characteristics include the size of the fund, the size of the family the fund belongs to, past performance, the age of the fund, the expense ratio, and the flows it received. I take the natural logarithm of the fund size, the fund family size, and the fund age. Among these control variables, the most important one is the fund family size because a large fund family is more likely to have a centralized AI or data science team. Furthermore, family size is also positively related to performance, as documented in Pástor et al. (2015). It is also important to include time fixed effects in the control variables, as the AI ratio is increasing over time by construction. A detailed definition of these variables is reported in Appendix A.

I implement the following Fama-MacBeth regressions:

$$Alpha_{i,t}^{Carhart} = \alpha + \beta\, AI_ratio_{i,t-2} + \gamma\, Controls_{i,t-1} + \eta_t + \delta_i + \epsilon_{i,t} \tag{9}$$

where the dependent variable is fund i's quarterly Carhart alpha. The AI ratio is lagged for one more period since it takes time for AI labor recruitment to affect fund performance. The regression is conducted at the quarterly level since the AI ratio and control variables are updated quarterly. Standard errors are clustered at the fund family and quarter level.

The regression results are reported in Table 3. Columns (1) to (4) correspond to different regression specifications. Consistent with the portfolio sorts, the results reveal a strong, statistically significant positive relationship between quarterly fund abnormal returns in quarter $q+1$ and the AI ratio in quarter $q-1$ across all regression specifications. Additionally, the results are economically meaningful. Given that the standard deviation of the AI ratio is 1.554% (see Table 1), the coefficient in Column (4) suggests that a 1-standard-deviation

higher AI ratio is associated with an annualized 61.4 basis points higher return ($0.0987 \times 1.554 \times 4 = 0.614$).

It is worth noting that the regression specifications in Column (3) and Column (4) include fund fixed effects, while Column (1) and Column (2) do not. According to Pástor et al. (2017), with fund fixed effects, the coefficient beta is a weighted average, across funds, of the slope estimates from fund-by-fund time-series regressions. The regression results in Table 3 show that the AI-performance relationship is stronger in the time series than in the cross section. In other words, as mutual funds gradually adopt AI technology, their performance improves.

Overall, the results reported in this subsection and the previous subsection show that the AI measure I constructed can predict future mutual fund performance, and this predictability persists even after adding control variables.

3.4 Time-Varying Predictability

The results computed so far average across all quarters in my dataset. However, the adoption of AI technology in the mutual fund industry does not increase linearly. As reported in Figure 2, AI labor recruitment was relatively low before 2016Q4 and increased rapidly thereafter. In this subsection, I equally separate the whole sample period into two parts: the first half spans from 2011 to 2016, and the second half from 2017 to 2022. I then repeat the portfolio sorting and Fama-MacBeth regressions in these two subsamples to determine where the predictive power comes from.

There are two competing hypotheses. On one hand, because AI became more powerful after 2016, we would expect return predictability to be stronger in the second half. On the other hand, AI technology might exhibit a "first mover advantage," where pioneer AI adopters earn more profit initially. When all funds adopt AI technology, none may generate excess returns anymore. In this case, we would expect the return predictability to be stronger in the first half.

Table 4 reports the results of a portfolio sorting analysis for different subsamples. The structure of the table is the same as Table 2. The performance differences in return (alpha) between the portfolios of high-AI and low-AI funds become more significant during 2017 to 2022.[15] The monthly performance difference becomes 0.241% for raw returns, 0.262% for CAPM alphas, and 0.094% for Carhart 4-factor alphas. Meanwhile, the performance differences are not significant during 2011 to 2016. Table 5 reports the Fama-MacBeth regression for different subsamples. The structure of the table is the same as Table 3. Consistent with the portfolio sorts, the AI ratio is significantly positively correlated with future alpha during 2017 to 2022 but cannot predict future alpha during 2011 to 2016. Panel B and Panel C in Table 6 report the factor loadings for different subsamples. The pattern in SMB loadings is mainly driven by the sample from 2017 to 2022. All these results support the first hypothesis that AI is more powerful in generating excess returns after 2016Q4. Given that AI technology is still rapidly evolving, mutual funds continue to benefit from its use, as we have not yet reached a steady state. The results in Figure 3 further support this, showing significant variation in AI usage among mutual fund companies even in 2022Q4. The lengthy process of developing a mature AI team means that not all mutual funds are yet leveraging AI to maximize profits.

4 Channel

Having established that the AI ratio can predict the future performance of mutual funds, a natural follow-up question concerns the underlying mechanism. In Section 2.1, I discussed several ways mutual funds might take advantage of AI technology. In this section, I test the hypothesis that mutual funds benefit from AI technology adoption by improving their information capacity. It is worth noting that all the potential channels are not mutually exclusive. I provide empirical evidence for one particular channel and do not exclude other

[15] The monthly raw returns are not significantly different from zero because the stock market fluctuated a lot during the COVID-19 period, leading to a high standard deviation.

potential channels (such as algorithmic trading) in my analysis.

4.1 Holding Analysis

It is well-documented that mutual fund managers face constraints on information processing due to limited cognitive resources and time. They can only allocate their attention to high-priority information, a phenomenon known as "rational inattention" in the literature (Kacperczyk et al., 2016; Ben-Rephael et al., 2017; Liu et al., 2023).[16] However, AI can serve as powerful tools for gathering and analyzing information, thereby enhancing mutual funds' information capacity. Moreover, AI is particularly adept at handling large, unstructured alternative data. Massa et al. (2024) study how institutional investors deal with big data and its impact on market efficiency. Sheng et al. (2024) also show that AI can help hedge fund managers process a large amount of unstructured data (e.g. conference call), which can aid managers' investment decisions.

To test whether mutual funds benefit from adopting AI technology by improving their information capacity, I focus on mutual funds' holdings. Cao et al. (2021) trained an AI analyst to predict stock returns using public information (e.g., corporate disclosures, macroeconomic indicators, etc.). They find that AI surpasses human analysts when the information is transparent but voluminous, while humans excel when critical information requires institutional knowledge or subjective judgment. My hypothesis is that if mutual funds increase their information capacity by adopting AI technology, they will tend to tilt their portfolios toward stocks with high information intensity, where they have a comparative advantage.[17]

To measure the information intensity of a stock, I adopt three measures from Cao et al.

[16]The concept "rational inattention" was first proposed by Sims (2003) in the field of macro-finance. Kacperczyk et al. (2016) was the first paper to introduce this concept to the mutual fund literature.

[17]Some may argue that mutual fund managers could simply trade in stocks where they have an informational advantage rather than holding them. However, mutual fund literature well documents that managers tend to hold these stocks. For example, a number of researchers study the geography of mutual fund investments. Coval and Moskowitz (1999) and Coval and Moskowitz (2001) find that fund managers exhibit a strong preference for investing in companies headquartered nearby due to their informational advantage. Additionally, Hong et al. (2005) and Pool et al. (2015) find that the holdings of fund managers who live in the same city/neighborhood are highly correlated because of social interactions.

(2021). The first measure is the number of information events, which refers to the number of firm-specific information events in Capital IQ Key Development data, representing the volume of available information about the firm. The second measure is firm size. Larger firms typically have more information available, whereas smaller firms often require more human subjective judgment. The third measure is firm age. The older a firm is, the more information tends to be available about it. I aggregate the stock-level measures to the fund level by taking the value-weighted average across fund holdings:

$$Holding_Information_Intensity_{i,t} = \sum_{k=1}^{N} w_{i,k,t} \times charateristic_{i,k,t} \tag{10}$$

where $w_{i,k,t}$ refers to the value of stock k held by fund i at quarter t divided by the total value of stocks held by fund i at quarter t. The term $charateristic_{i,k,t}$ refers to the three measures of stock-level information intensity. I test whether funds adopting AI technology hold stocks in which AI has a comparative advantage by estimating the following regression of stock information intensity on the AI ratio:

$$Holding_Information_Intensity_{i,t} = \alpha + \beta\, AI_ratio_{i,t-1} + \gamma\, Controls_{i,t-1} + \eta_t + \delta_i + \epsilon_{i,t} \tag{11}$$

where $Holding_Information_Intensity_{i,t}$ is the weighted average of the three measures of stock-level information intensity, as calculated in the previous equation. I control the fund fixed effect and time fixed effect in this regression.

Table 7 reports the regression outcomes. Columns (1) and (2) indicate that the average number of information events of stocks held by funds is approximately 0.7 (0.488 (0.421) $\times$ 1.554) higher for a 1-standard-deviation increase in the AI ratio, given that the mean of the AI ratio is 1.554%. The coefficient is significant at 1% level. Columns (3) and (4) indicate that the average market capitalization of stocks held by funds is approximately 26 million dollar (17.894 (16.268) $\times$ 1.554) higher for a 1-standard-deviation increase in the AI ratio. Columns (5) and (6) indicate that the average age of stocks held by funds is approximately

0.22 (0.149 (0.132) $\times$ 1.554) higher for a 1-standard-deviation increase in the AI ratio. Taken together, mutual funds with higher AI ratio tend to tend to tilt their portfolios toward large stock, old stock and stock with more information events. Overall, these results suggest that mutual funds adopting AI technology outperform other funds by holding high information intensity stocks, where they have a comparative advantage.

4.2 Identification

In this subsection, I provide additional evidence that mutual funds with a higher AI ratio tend to hold high information intensity stocks by utilizing a shock in AI technology. My hypothesis is that when AI technology improves, mutual funds with a high AI labor stock will hold higher information intensity stocks, whereas funds with low AI labor stock will not.

For the AI technology shock, I rely on the publication of the Transformer architecture in June 2017.[18] The Transformer is a deep learning architecture developed by Google, based on the multi-head attention mechanism, and was proposed in the paper "Attention Is All You Need". This paper has become one of the most highly cited in the field of AI (128,482 citations by August 2024). Since its publication, the Transformer architecture has become a foundational architecture in various areas of AI, including machine learning, natural language processing, and image recognition. Most pre-trained models are based on Transformer architecture, such as Clip and GPT.

I run difference-in-difference regressions to test the hypothesis. The treatment group consists of funds whose AI ratio is above the median, while the control group consists of funds whose AI ratio is below the median in June 2016.[19] I include fund fixed effects and time fixed effects in the regression as control variables. The sample period spans from 2016Q3

[18] Another breakthrough in AI technology is ChatGPT, launched on November 30, 2022. I do not use it because my sample ends in December 2022.

[19] The result remains robust if I change the formation time, as the AI ratio has an average autocorrelation higher than 90%.

to 2018Q2, covering one year before and after the cutoff. The regression is specified as follows:

$$Holding_Information_Intensity_{i,t} = \alpha + \beta\,(Post_{t-1} \times Treatment_i) + \gamma\, Controls_{i,t-1} + \eta_t + \delta_i + \epsilon_{i,t} \tag{12}$$

where the dependent variable $Holding_Information_Intensity_{i,t}$ is the same as before. $Treatment_i$ equals to one if fund i's AI ratio is above the median in June 2016. $Post_t$ equals one from 2017Q2 onward (in other words, $Post_{t-1}$ equals one from 2017Q3). All the independent variables and control variables are lagged by one quarter, as it takes time for the AI technology shock to become effective. I verify the parallel assumption by plotting the average dependent variables of the treatment group and control group. Figure 4 shows that there is almost no pre-trend for the different dependent variables before 2017Q2.

Table 8 reports the results of the difference-in-difference regressions. The results show that following the breakthrough in AI technology, mutual funds with a higher AI labor stock tend to shift their portfolio allocations toward larger stocks and those with more information events. In the four quarters after the publication of the Transformer model, the average number of information events and market capitalization of stocks held by the treatment group are 0.9 and 6.6 million dollars higher, respectively, compared to those held by the control group. Another dependent variable, the age of the stock, becomes insignificant in this setting. These results suggest that mutual funds with higher AI labor stock tilt their portfolios toward stocks with high information intensity after the AI technology shock.

4.3 Future Performance of Stock Purchased/Sold by AI Fund

The previous two subsections provide evidence that mutual funds tend to hold stocks with higher information intensity after adopting AI technology. I investigate whether the outperformance of these AI-enhanced mutual funds is driven by trading in stocks with higher information intensity. My methodology largely follows Bai et al. (2023). If the outperformance is indeed attributable to high information intensity stocks, I expect trades in these

stocks to be more profitable than those in stocks with lower information intensity.

First, I sort all mutual funds into quintiles based on their AI ratio, following the method outlined in Section 3.2. I then aggregate the holdings of all mutual funds within each quintile into a single portfolio. Next, I calculate the change in portfolio weight resulting from active rebalancing, categorizing these changes into buy and sell trades. Specifically, I begin by computing the portfolio's hypothetical weights for a given quarter, assuming no trading occurs, denoted as $\hat{w}_{k,t}$.

$$\hat{w}_{k,t} = \frac{w_{k,t-1}(1 + r_{k,t})}{\sum_{k=1}^{N} w_{k,t-1}(1 + r_{k,t})} \tag{13}$$

where $w_{k,t-1}$ is portfolio Q's weight in stock k at the end of quarter $t-1$ and $r_{k,t}$ is stock k's return in quarter t. My calculation for $\hat{w}_{k,t}$ indicates that if no trades occur in quarter t, changes in portfolio weights are solely driven by stock returns during that period. Next, I determine "buy" and "sell" trades by comparing the hypothetical portfolio weights with the active portfolio weights. Specifically, $\hat{w}_{k,t} < w_{k,t}$ suggests a "buy" trade, increasing the weight associated with stock k, while $\hat{w}_{k,t} > w_{k,t}$ indicates a "sell" trade in stock k. I separate the stocks into two categories (high/low information intensity) based on the median number of information events each quarter. I use the characteristic-selectivity measure developed by Daniel et al. (1997), hereafter referred to as DGTW, to assess stock returns across different characteristics. Finally, I calculate the future performance of buy versus sell trades within each AI quintile portfolio. Specifically, for each portfolio, I compute the mean quarterly DGTW benchmark-adjusted returns of stocks that were purchased or sold during the preceding quarter.

My findings reveal a significant difference in stock performance between buy and sell trades among funds with a high AI ratio, especially when trading stocks with high information intensity. Table 9 presents the results. In the top quintile (Q5), the average quarterly DGTW-adjusted return is 0.62% for stocks purchased with high information intensity, compared to -0.49% for stocks sold with high information intensity. This indicates that superior trade

performance for funds with a high AI ratio is primarily driven by their buy trades. Moreover, this performance is enhanced by trades in high information intensity stocks. For example, in Column (3), the difference between buy and sell trades is 1.11% (t-stat=2.49) for stocks with high information intensity, while it becomes -0.23% (t-stat=-0.91) for stocks with low information intensity. Overall, this evidence supports that the outperformance of mutual funds adopting AI technology stems from trading stocks with high information intensity.

4.4 Manager Skill: Stock Picking and Market Timing

The results in the previous subsections show that mutual funds improve their information capacity after adopting AI technology, evidenced by the information intensity of their stock holdings. In this subsection, I investigate the source of outperformance for these AI-enhanced mutual funds using another method. Kacperczyk et al. (2014) evaluate mutual fund manager skill by decomposing fund performance into two parts: stock picking and market timing. If adopting AI technology provides mutual funds with an advantage in analyzing stocks with high information intensity, I expect that the outperformance should come primarily from improved stock picking skills.

I calculate the stock picking skill and market timing skill in fund-quarter level following the method of Kacperczyk et al. (2014). Market timing skill refers to a manager's ability to adjust their portfolio's exposure to the market according to prevailing market conditions. $Timing_{i,t}$ denotes the market timing skill for fund i at quarter t, calculated as follows:

$$Timing_{i,t} = \sum_{k=1}^{N_i} (w_{i,k,t} - w^m_{k,t})(\beta_{k,t} R^m_{t+1}) \tag{14}$$

where β_k is the covariance of stock k's return, R_k, with the market return, R^m, divided by the variance of the market return. It is calculated using a 60 months moving window (with at least 24 months of nonmissing return). The portfolio weight, denoted as $w_{i,k,t}$, is the fraction of fund i's total asset held in stock k at the start of quarter t. The market weight,

denoted as $w^m_{k,t}$, is the fraction of total market capitalization in stock k. Similarly, stock picking skill refers to a manager's ability to adjust holdings of individual stocks based on their idiosyncratic returns relative to the broader market. $Picking_{i,t}$ denotes stock picking skill for fund i at quarter t, calculated as follows:

$$Picking_{i,t} = \sum_{k=1}^{N_i} (w_{i,k,t} - w^m_{k,t})(R_{k,t+1} - \beta_{k,t} R^m_{t+1}) \tag{15}$$

A fund with a high stock picking ability overweighs stocks that subsequently have high idiosyncratic returns and underweights those with low idiosyncratic returns. The summary statistic of these two skills are also reported in Panel A of Table 1.

Similar to Equation (11), I test whether funds adopting AI technology have a higher stock picking skill or market timing skill by estimating the following regression of skills on the AI ratio:

$$Skill_{i,t} = \alpha + \beta\, AI_ratio_{i,t-1} + \gamma\, Controls_{i,t-1} + \eta_t + \delta_i + \epsilon_{i,t} \tag{16}$$

where $Skill_{i,t}$ refers to the stock picking skill or market timing skill of fund i at quarter t, as calculated in the previous equations. I also control the fund fixed effect and time fixed effect in this regression.[20]

Table 10 reports the regression outcomes. The results show that mutual funds adopting AI technology exhibit higher skill in stock picking. Given that the standard deviation of the AI ratio is 1.554% (see Table 1), the coefficient in Column (2) suggests that a 1-standard-deviation higher AI ratio is associated with a 46.6 basis point annual advantage in stock picking ($0.075 \times 1.554 \times 4 = 0.466$). The coefficients in Column (3) and Column (4) are also positive but not significant, indicating that mutual funds might slightly improve their market timing skill by adopting AI technology. All these results are consistent with the hypothesis that mutual funds benefit from AI technology adoption by improving their information

[20] Kacperczyk et al. (2016) find that mutual funds allocate more attention on stock picking in booms and market timing in recessions. I control this effect by adding time fixed effect into the regression.

capacity.

5 Impact on Mutual Fund Managers

Another important research question in the AI area is to study its impact on the labor market, as many people fear that AI will take their jobs. For example, Acemoglu et al. (2022) examines the impact of AI-labor substitution on employment and wage growth. In this section, I explore how AI technology adoption affects mutual fund managers. To answer this question, I investigate the relationship between the AI ratio and manager turnover at mutual funds. On one hand, if mutual funds rely more on AI technology, they may reduce their dependence on individual fund managers. On the other hand, AI technology may not easily threaten mutual fund managers, as this is a high-tech occupation that requires numerous soft skills. Therefore, whether the AI ratio can predict higher fund manager turnover remains an open question.

I construct two manager turnover variables as dependent variables, following Kostovetsky and Warner (2015).[21] The first variable is a manager turnover dummy, which takes a value of one if a manager departs (and the fund survives) in a given quarter and zero otherwise. The second variable, manager turnover, is an adjustment of the manager turnover dummy based on the total number of managers. For instance, if two out of five managers leave, the manager turnover equals 0.4 for that quarter. To examine the relationship between manager turnover and AI ratio, I employ both OLS and probit regressions, with manager turnover (dummy) as the dependent variable and AI ratio as the main independent variable. Following Kostovetsky and Warner (2015), I control for several variables that can affect manager turnover, such as team size and past performance (measured by alpha over the past year). The AI ratio is also lagged by one more quarter.

Table 11 reports the regression outcomes. I find a negative relationship between manager

[21] The historical mutual fund managers list is obtained from Morningstar Direct and linked to CRSP using CUSIP.

turnover and fund size, and a positive relationship between manager turnover and fund family size/team size. All these findings are consistent with Kostovetsky and Warner (2015). However, the main independent variable, the AI ratio, is insignificant in all the regression specifications, indicating that the AI ratio cannot predict manager turnover. These results suggest that although AI technology is powerful, it has not yet threatened the positions of mutual fund managers.

6 Conclusion

In this paper, I study how AI technology, one of the most important new technologies of the last decade, shapes the mutual fund industry. I develop a new measure of AI technology adoption for mutual funds, derived from AI labor recruitment data based on job postings from Burning Glass Technologies. This unique measure allows me to test whether the adoption of AI technology can predict mutual fund performance. I find that a long-short portfolio, which goes long in the top quintile of funds with the highest AI ratio and short in the bottom quintile of funds with the lowest AI ratio, delivers an annual excess return of 156 basis points. This predictability primarily comes from the second half of the sample period.

I also study the underlying mechanism of the return predictability. I test the hypothesis that mutual funds benefit from adopting AI technology by improving their information capacity. I provide empirical evidence for this channel by showing that mutual funds with a high AI ratio tilt their portfolios toward stocks with high information intensity, where they have a comparative advantage. These results suggests that the large-scale use of AI by institutional investors may also affect the price informativeness of the stock market. I leave this question for future research.

Reference

Abis, S. (2020): "Man vs. machine: Quantitative and discretionary equity management," *Available at SSRN 3717371*.

Abis, S. and L. Veldkamp (2024): "The changing economics of knowledge production," *The Review of Financial Studies*, 37, 89–118.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022): "Artificial intelligence and jobs: Evidence from online vacancies," *Journal of Labor Economics*, 40, S293–S340.

Avramov, D., S. Cheng, and L. Metzker (2023): "Machine learning vs. economic restrictions: Evidence from stock return predictability," *Management Science*, 69, 2587–2619.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024): "Artificial intelligence, firm growth, and product innovation," *Journal of Financial Economics*, 151, 103745.

Bai, J. J., Y. Tang, C. Wan, and H. Z. Yuksel (2023): "Thematic investing in mutual funds," *Zafer, Thematic Investing in Mutual Funds (September 19, 2023). Northeastern U. DAmore-McKim School of Business Research Paper*.

Ben-Rephael, A., Z. Da, and R. D. Israelsen (2017): "It depends on where you search: Institutional investor attention and underreaction to news," *The Review of Financial Studies*, 30, 3009–3047.

Bertomeu, J., Y. Lin, Y. Liu, and Z. Ni (2023): "Capital market consequences of generative AI: Early evidence from the ban of ChatGPT in Italy," *Available at SSRN*, 4452670.

Bonelli, M. (2023): "Data-driven Investors," *Available at SSRN 4362173*.

Bonelli, M. and T. Foucault (2023): "Displaced by Big Data: Evidence from Active Fund Managers," *Available at SSRN 4527672*.

Cao, S., Y. Cheng, M. Wang, Y. Xia, and B. Yang (2022): "Visual Information in the Age of AI: Evidence from Corporate Executive Presentations," *Available at SSRN 4490834*.

Cao, S., W. Jiang, J. L. Wang, and B. Yang (2021): "From man vs. machine to man+ machine: The art and AI of stock analyses," *Columbia Business School Research Paper*.

Chen, L., M. Pelger, and J. Zhu (2024): "Deep learning in asset pricing," *Management Science*, 70, 714–750.

Chen, R. and J. Ren (2022): "Do AI-powered mutual funds perform better?" *Finance Research Letters*, 47, 102616.

Chen, Y., B. T. Kelly, and D. Xiu (2022): "Expected returns and large language models," *Available at SSRN 4416687*.

Coval, J. D. and T. J. Moskowitz (1999): "Home bias at home: Local equity preference in domestic portfolios," *The Journal of Finance*, 54, 2045–2073.

——— (2001): "The geography of investment: Informed trading and asset prices," *Journal of political Economy*, 109, 811–841.

DAcunto, F., N. Prabhala, and A. G. Rossi (2019): "The promises and pitfalls of robo-advising," *The Review of Financial Studies*, 32, 1983–2020.

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

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

Doshi, H., R. Elkamhi, and M. Simutin (2015): "Managerial activeness and mutual fund performance," *The Review of Asset Pricing Studies*, 5, 156–184.

Dou, W. W., I. Goldstein, and Y. Ji (2024): "Ai-powered trading, algorithmic collusion, and price efficiency," *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.

Du, K., M. Liu, and S. Wang (2023): "Human Information Production in the Machine Age: Evidence from Automated Information Acquisition in the Asset Management Industry," *Available at SSRN 4408892*.

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

Gabaix, X., R. S. Koijen, R. Richmond, and M. Yogo (2023): "Asset embeddings," *Available at SSRN 4507511*.

Gofman, M. and Z. Jin (2024): "Artificial intelligence, education, and entrepreneurship," *The Journal of Finance*, 79, 631–667.

Grennan, J. and R. Michaely (2020): "Artificial intelligence and high-skilled work: Evidence from analysts," *Swiss Finance Institute Research Paper*.

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

Hong, H., J. D. Kubik, and J. C. Stein (2005): "Thy neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers," *The Journal of Finance*, 60, 2801–2824.

Kacperczyk, M., S. V. Nieuwerburgh, and L. Veldkamp (2014): "Time-varying fund manager skill," *The Journal of Finance*, 69, 1455–1484.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp (2016): "A rational theory of mutual funds' attention allocation," *Econometrica*, 84, 571–626.

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

Kim, A., M. Muhn, and V. V. Nikolaev (2024): "Bloated disclosures: can ChatGPT help investors process information?" *Chicago Booth Research Paper*, 2023–59.

Kostovetsky, L. and J. B. Warner (2015): "Youre fired! New evidence on portfolio manager turnover and performance," *Journal of Financial and Quantitative Analysis*, 50, 729–755.

Leonardo Gambacorta, F. S. and S. Schiaffi (2024): "Artificial intelligence and relationship lending," *Working Paper*.

Li, B., A. Rossi, S. Yan, and L. Zheng (2022): "Real-time machine learning in the cross-section of stock returns," Tech. rep., Working Paper.

Li, B. and A. G. Rossi (2020): "Selecting mutual funds from the stocks they hold: A machine learning approach," *Available at SSRN 3737667*.

Liu, H., L. Peng, and Y. Tang (2023): "Retail attention, institutional attention," *Journal of Financial and Quantitative Analysis*, 58, 1005–1038.

Lopez-Lira, A. and Y. Tang (2023): "Can chatgpt forecast stock price movements? return predictability and large language models," *arXiv preprint arXiv:2304.07619*.

Lou, D. (2012): "A flow-based explanation for return predictability," *The Review of Financial Studies*, 25, 3457–3489.

Lu, F., L. Huang, and S. Li (2023): "ChatGPT, Generative AI, and Investment Advisory," *Available at SSRN 4519182*.

Massa, M., H. Zhang, and Y. Zhou (2024): "Data Specialists and Market Efficiency," *Available at SSRN 4739691*.

Pástor, L., R. F. Stambaugh, and L. A. Taylor (2015): "Scale and skill in active management," *Journal of Financial Economics*, 116, 23–45.

——— (2017): "Do funds make more when they trade more?" *The Journal of Finance*, 72, 1483–1528.

Pool, V. K., N. Stoffman, and S. E. Yonker (2015): "The people in your neighborhood: Social interactions and mutual fund portfolios," *The Journal of Finance*, 70, 2679–2732.

Sheng, J., S. Zheng, B. Yang, and A. Zhang (2024): "Generative AI and Asset Management," *Available at SSRN 4786575*.

Sheng-Syan Chen, J. K. and S.-C. Peng (2024): "Harnessing Artificial Intelligence: Evidence from Corporate Investment Outcomes and Efficiency," *Working Paper*.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50, 665–690.

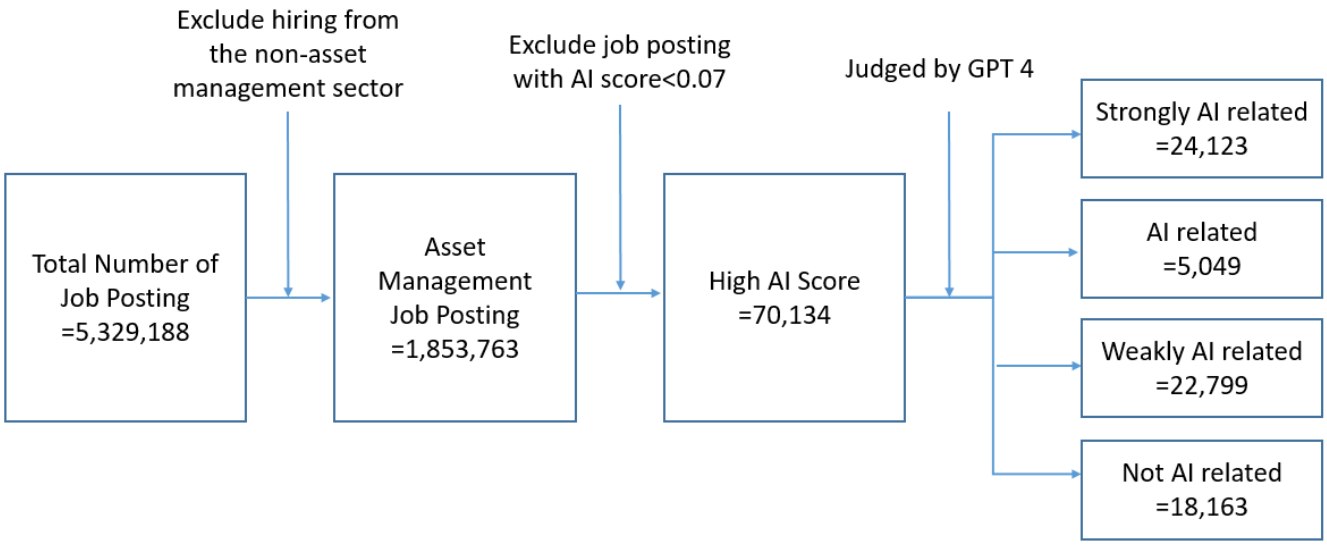


Figure 1: **Data Cleaning Process** This figure shows the process to clean the Burning Glass data and indentify AI jobs. It also shows the number of observations in each step. The sample period is from 2010 to 2022.

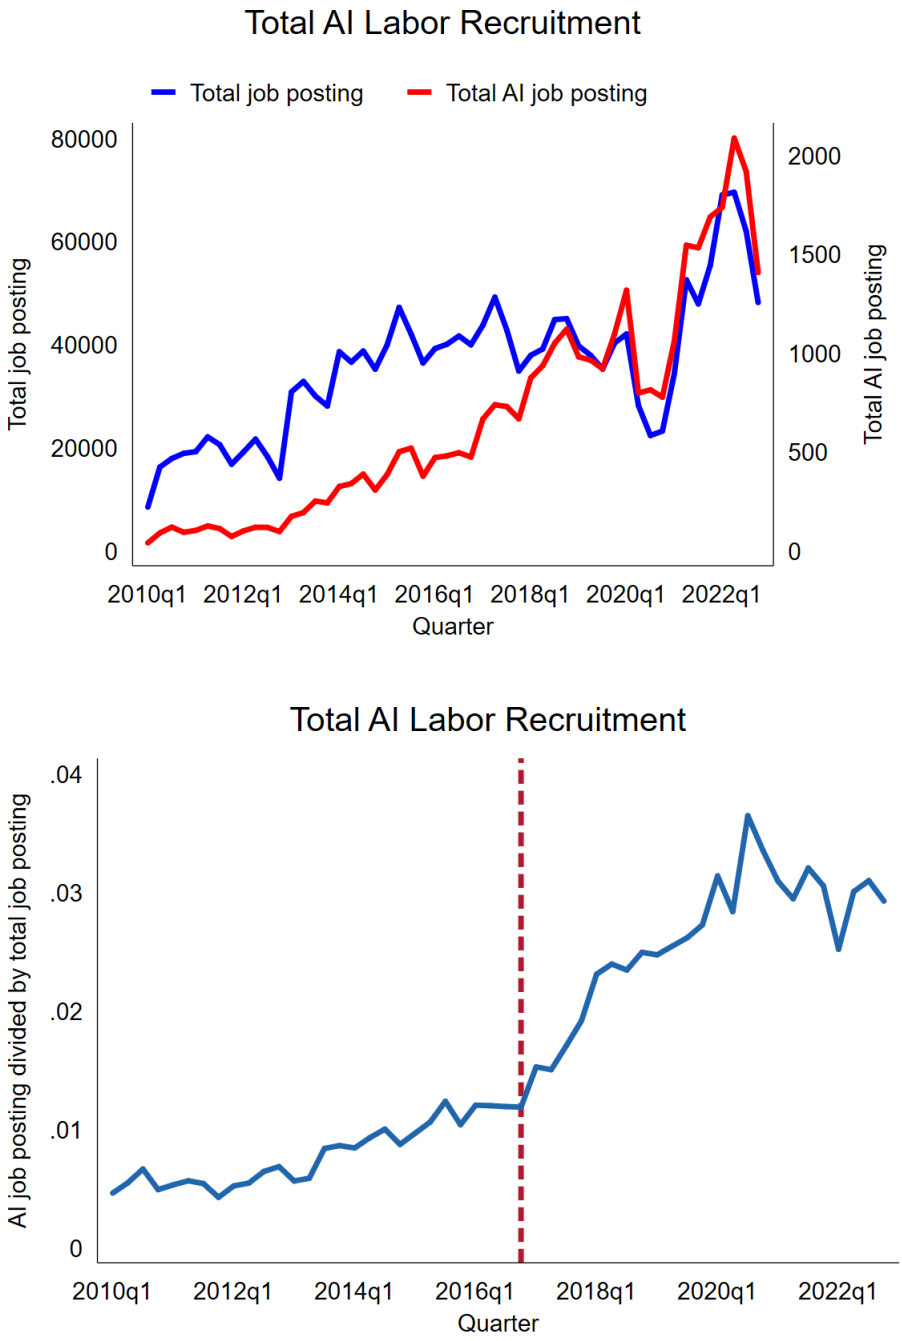


Figure 2: **AI Labor Recruitment** These figures plot the AI labor recruitment in the mutual fund industry quarterly. In the first figure, the red line is the total number of AI job posting for all the fund companies, correspond to the right y-axis. The blue line is the total number of all the job posting, correspond to the left y-axis. In the second figure, the y-axis is the total number of AI job posting divided by total number of job posting. The dash line stands for 2016Q4.

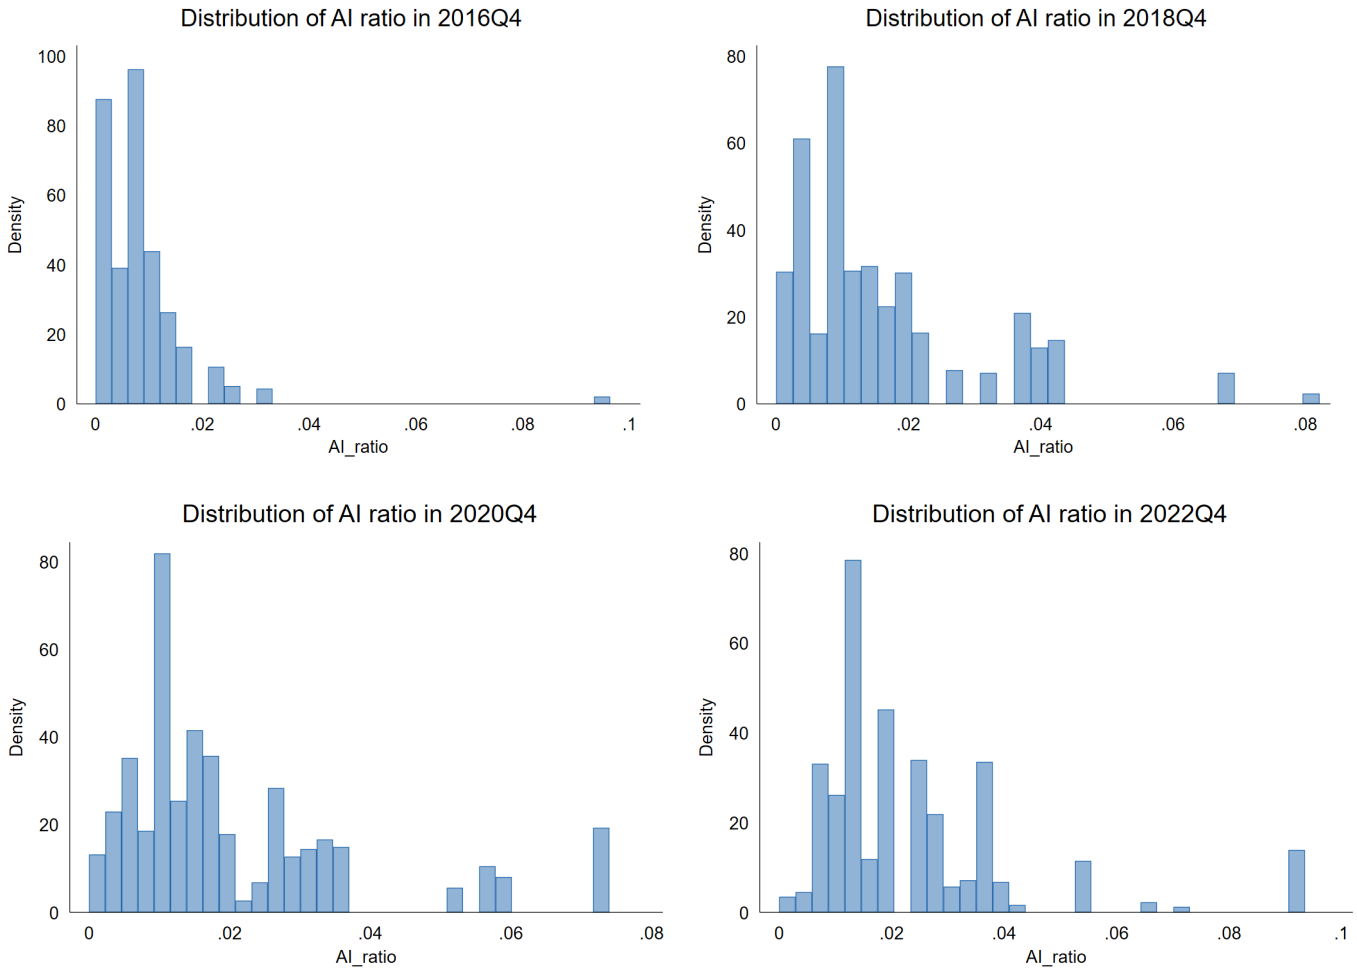


Figure 3: **AI Ratio Distribution**. These figures plot the distribution of AI ratio across all mutual funds at different points in time. The x-axis represents the AI ratio, while the y-axis indicates the number of funds in each bin. The four subfigures correspond to the distributions in 2016Q4, 2018Q4, 2020Q4, and 2022Q4, respectively.

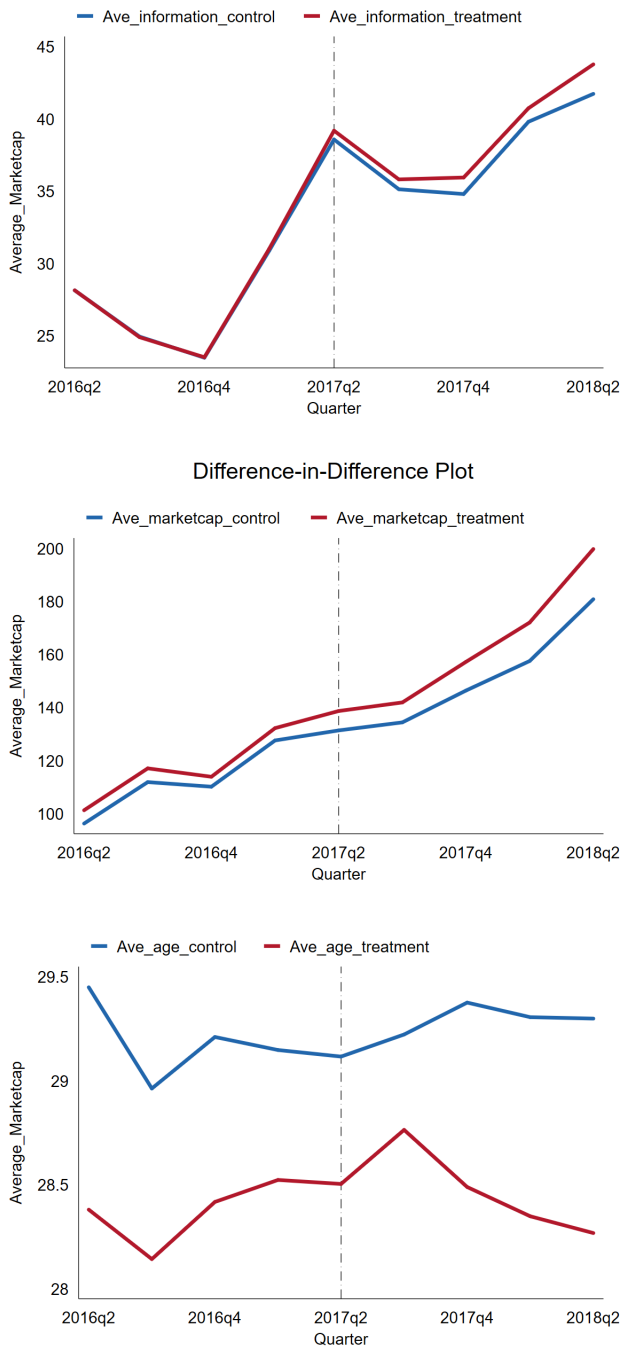


Figure 4: **Difference-in-Difference Plot** These figures plot the average dependent variables of the treatment group and control group in the difference-in-difference regressions (e.g., Equation 12). The x-axis represents time. In the top panel, the y-axis represents the number of information events. In the middle panel, the y-axis represents the market cap. In the bottom panel, the y-axis represents the firm age. The dotted vertical line indicates the event time, the second quarter of 2017, when the Transformer was proposed.

Table 1: **Summary Statistic**

This table reports the summary statistic of fund characteristic. Panel A reports the number of observation, the mean, the standard deviation and different percentiles. Panel B reports the correlation matrix of the fund characteristic. A full description of all the variables can be found in the appendix.

Panel A: Summary Statistic

	N	Mean	Std	p1	p25	p50	p75	p99
AI ratio (%)	41,121	1.369	1.554	0	0.424	0.889	1.723	8.214
TNA (in millions)	41,121	4,047	11,207	17.80	265	834.6	2,584	82,383
Family size	40,112	279,330	467,368	198	20,789	42,621	383,393	2,254,000
Qret (%)	40,974	2.640	8.991	-25.23	-0.563	3.452	7.515	26.13
Flow	40,848	-0.0009	0.106	-0.285	-0.037	-0.0164	0.0124	0.599
Num holdings	41,014	240.1	456.7	1	53	92	201	2,570
Alpha (%)	38,542	-0.119	2.648	-9.399	-1.357	-0.00745	1.236	7.759
Turnover (%)	30,332	0.562	0.449	0.0300	0.240	0.460	0.760	2.340
Expenses (%)	41,109	0.568	0.469	0	0	0.659	0.960	1.560
Age (in years)	41,110	17.20	11.89	1	9.005	14.26	22.52	56.79
Activeshare	32,048	0.854	0.158	0.336	0.787	0.904	0.977	0.999
Holding Marketcap	37,050	190.4	252.0	1.332	9.895	101.9	252.9	1,114
Holding Information	37,050	33.08	18.91	9.455	14.74	32.26	47.09	83.15
Holding Age	37,050	29.25	9.772	11.21	21.94	27.70	35.06	56.46
Timing (%)	37,359	2.497	7.506	-20.57	-0.010	3.202	6.377	20.18
Picking (%)	37,359	-0.0425	3.718	-10.03	-1.641	0	1.468	9.876
Manager Turnover	38,635	0.0225	0.114	0	0	0	0	0.500

Panel B: Correlation

	AI ratio	TNA	Age	Qret	Flow	Turnover	Expenses	Activeshare
AI ratio (%)	1							
TNA (in millions)	0.138	1						
Age (in years)	0.0301	0.329	1					
Qret	-0.0100	0.0168	0.00177	1				
Flow	-0.00544	0.00378	-0.132	0.0369	1			
Turnover (%)	-0.0304	-0.222	-0.0415	0.00781	-0.0293	1		
Expenses	-0.169	-0.110	0.312	0.00695	-0.0560	0.347	1	
Activeshare	-0.0617	-0.295	-0.0325	-0.0178	-0.0171	0.240	0.265	1

Table 2: **Portfolio Sorting**

This table reports the results of a portfolio sorting analysis. Mutual funds are sorted into 5 portfolios based on their AI measure at the beginning of each semi-year. I calculate the average performance for each portfolio each month, value-weighted by TNA. The three columns report the time-series averages of raw returns, CAPM Alpha, and Carhart Alpha, respectively. All performance measures are expressed as percentages per month. The bottom row reports the mean monthly return (alpha) differences between the portfolios of high-AI (top quintile) and low-AI (bottom quintile) funds. t-Statistics are provided in parentheses.

Quintile	Raw Return	CAPM Alpha	Carhart Alpha
1 (Low)	0.822**	-0.180***	-0.065
	(2.233)	(-3.540)	(-1.585)
2	0.876**	-0.075**	-0.009
	(2.455)	(-2.136)	(-0.272)
3	0.931***	-0.056	-0.031
	(2.655)	(-1.250)	(-0.708)
4	0.872**	-0.097**	-0.034
	(2.456)	(-2.059)	(-0.734)
5 (High)	0.952***	-0.016	0.011
	(2.684)	(-0.446)	(0.348)
Difference: High-low	0.130***	0.164***	0.076**
	(3.127)	(4.202)	(2.406)
Observation		144	

Table 3: **AI and Future Performance**

This table reports the results of the Fama-MacBeth regression of quarterly Carhart alphas (in percentage) in quarter q+1 on fund characteristics measured at the end of quarter q and AI ratio measured at the end of quarter q-1 (Equation 9 in Section 3.3). Different columns include various control variables and fixed effects. The "Alpha" in the second row refers to the Carhart alpha from the previous quarter, included as a control variable. Standard errors are clustered at the fund family and quarter levels; t-statistics are reported in parentheses.

Carhart Alpha	(1)	(2)	(3)	(4)
AI ratio (%)	0.0663***	0.0603***	0.0813**	0.0987**
	(3.13)	(3.12)	(2.56)	(2.16)
Alpha (%)		0.0083		-0.0353
		(0.22)		(-0.94)
Logsize		-0.0074		-0.3645***
		(-0.51)		(-3.30)
Logage		-0.0751**		0.1631
		(-2.07)		(0.56)
Flow (%)		-1.2317***		-1.2909***
		(-6.79)		(-6.39)
Expenses (%)		-0.097		-0.381**
		(-1.42)		(-2.63)
Logfamilysize		0.0125		0.0632
		(1.46)		(1.13)
Observations	35,559	34,173	35,546	34,159
R-squared	0.104	0.108	0.143	0.150
Fund FE	NO	NO	YES	YES
Time FE	YES	YES	YES	YES

Table 4: **Portfolio Sorting: Different Subsample**

These tables report the results of a portfolio sorting analysis for different subsample. Panel A reports the results in sample period from 2011 to 2016. Panel B reports the results in sample period from 2017 to 2022. Mutual funds are sorted into 5 portfolios based on their AI measure at the beginning of each semi-year. I calculate the average performance for each portfolio each month, value-weighted by TNA. The three columns report the time-series averages of raw returns, CAPM Alpha, and Carhart Alpha, respectively. All performance measures are expressed as percentages per month. The bottom row reports the mean monthly return (alpha) differences between the portfolios of high-AI (top quintile) and low-AI (bottom quintile) funds. t-Statistics are provided in parentheses.

Panel A: Sample Period 2011-2016			
Quintile	Raw Return	CAPM Alpha	Carhart Alpha
1 (Low)	0.897**	-0.145***	-0.088*
	(2.182)	(-3.010)	(-1.808)
2	0.898**	-0.122**	0.008
	(2.193)	(-2.217)	(0.164)
3	1.001**	-0.049	0.004
	(2.488)	(-0.809)	(0.071)
4	0.934**	-0.109**	0.012
	(2.259)	(-2.100)	(0.254)
5 (High)	0.920**	-0.075**	-0.029
	(2.266)	(-2.560)	(-0.903)
Difference: High-low	0.024	0.070	0.059
	(0.519)	(1.550)	(1.356)
Observation		72	
Panel B: Sample Period 2017-2022			
Quintile	Raw Return	CAPM Alpha	Carhart Alpha
1 (Low)	0.745	-0.216**	-0.041
	(1.198)	(-2.374)	(-0.613)
2	0.853	-0.026	-0.028
	(1.435)	(-0.615)	(-0.593)
3	0.857	-0.062	-0.068
	(1.469)	(-0.950)	(-0.951)
4	0.807	-0.084	-0.083
	(1.376)	(-1.055)	(-1.019)
5 (High)	0.986	0.045	0.053
	(1.666)	(0.683)	(0.963)
Difference: High-low	0.241***	0.262***	0.094**
	(3.533)	(4.196)	(2.029)
Observation		72	

Table 5: **AI and Future Performance: Different Subsample**

These tables report the results of the Fama-MacBeth regression of quarterly Carhart alphas (in percentage) in quarter q+1 on fund characteristics measured at the end of quarter q and AI ratio measured at the end of quarter q-1. Panel A reports the results in sample period from 2011 to 2016. Panel B reports the results in sample period from 2017 to 2022. Different columns include various control variables and fixed effects. Standard errors are clustered at the fund family and quarter levels; t-statistics are reported in parentheses.

Panel A: Sample Period 2011-2016				
Carhart Alpha	(1)	(2)	(3)	(4)
AI ratio (%)	0.0354	0.0275	0.0245	-0.0070
	(1.62)	(1.63)	(0.28)	(-0.07)
Alpha (%)		-0.0953		-0.1361**
		(-1.64)		(-2.24)
Logsize		-0.0372***		-0.9742***
		(-3.68)		(-3.59)
Logage		-0.0282		0.5703
		(-0.46)		(1.06)
Flow (%)		-1.5240***		-1.3693***
		(-5.95)		(-4.44)
Expenses (%)		-0.133		-0.655**
		(-1.59)		(-2.45)
Logfamilysize		0.0172		0.0998*
		(1.39)		(1.84)
Observations	14,636	14,127	14,617	14,107
R-squared	0.073	0.088	0.111	0.138
Panel B: Sample Period 2017-2022				
Carhart Alpha	(1)	(2)	(3)	(4)
AI ratio (%)	0.0737***	0.0625***	0.1175**	0.1314**
	(3.12)	(3.08)	(2.23)	(1.99)
Alpha (%)		0.0894*		0.0177
		(1.88)		(0.37)
Logsize		0.0154		-0.5150**
		(0.63)		(-2.42)
Logage		-0.1013***		0.2882
		(-2.97)		(0.49)
Flow (%)		-0.8288***		-0.7966**
		(-3.44)		(-2.60)
Expenses (%)		-0.092		-0.216
		(-0.91)		(-0.85)
Logfamilysize		0.0037		-0.1620
		(0.46)		(-0.80)
Observations	19,473	18,631	19,466	18,623
R-squared	0.129	0.139	0.188	0.199
Fund FE	NO	NO	YES	YES
Time FE	YES	YES	YES	YES

Table 6: **Portfolio Sorting: Factor Loading**

This table summarizes the factor loadings from Carhart four-factor model regressions using net-of-expenses returns. Mutual funds are sorted into 5 portfolios based on their AI measure at the beginning of each semi-year. The factor loadings are calculated using the daily returns for each fund each quarter. I calculate the average factor loading for each portfolio each quarter, value-weighted by TNA. Panel A reports the results for the full sample. Panel B reports the results for the sample period from 2011 to 2016. Panel C reports the results for the sample period from 2017 to 2022.

Panel A: Full Sample				
Group	MKT	SMB	HML	MOM
1	0.9503	0.0963	0.0474	-0.0101
2	0.9522	0.0488	-0.0269	0.0000
3	0.9570	0.0333	-0.0575	0.0145
4	0.9557	0.0260	-0.0581	0.0089
5	0.9627	0.0392	0.0183	0.0047
Panel B: Sample Period 2011-2016				
Group	MKT	SMB	HML	MOM
1	0.9410	0.0186	-0.0126	-0.0108
2	0.9573	0.0658	-0.0262	-0.0043
3	0.9702	0.0440	-0.0893	0.0190
4	0.9654	0.0172	-0.0430	0.0170
5	0.9662	0.0218	-0.0046	0.0077
Panel C: Sample Period 2017-2022				
Group	MKT	SMB	HML	MOM
1	0.9600	0.1774	0.1100	-0.0095
2	0.9469	0.0310	-0.0276	0.0046
3	0.9433	0.0220	-0.0243	0.0098
4	0.9456	0.0352	-0.0739	0.0005
5	0.9591	0.0574	0.0422	0.0016

Table 7: **Holding Analysis**

This table reports the impact of AI on mutual fund holding. I choose three stock level measures and aggregate to fund level by calculating the weighted average. The three measures are: number of information event, market capitalization and age of stock The independent variable is the AI ratio. All the independent variable and control variables are lagged for one quarter. A full description of all the variables can be found in the appendix. Standard errors are clustered at the fund family and quarter level; t-statistics are reported in parentheses.

	Information Event		Marketcap		Age of Stock	
	(1)	(2)	(3)	(4)	(5)	(6)
AI ratio (%)	0.488***	0.421***	17.894**	16.268**	0.149**	0.132*
	(2.89)	(2.75)	(2.05)	(2.05)	(2.04)	(1.99)
Qret (%)		0.061		0.978		-0.004
		(0.58)		(0.47)		(-0.25)
Logsize		1.149**		31.451***		-0.085
		(2.41)		(3.64)		(-0.84)
Logfamilysize		-0.912*		-8.946		0.309*
		(-1.86)		(-1.21)		(1.72)
Logage		-1.282**		-33.516**		0.045
		(-2.23)		(-2.35)		(0.12)
Flow		-2.362***		-26.128**		0.288
		(-3.72)		(-2.19)		(0.86)
Expenses (%)		-0.236		16.71		-0.285
		(-0.37)		(0.99)		(-0.59)
Observations	35,562	34,456	35,270	34,163	35,191	34,084
R-squared	0.840	0.841	0.773	0.779	0.912	0.913
Fund FE	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES

Table 8: **Identification: Difference in Difference**

This table reports the difference-in-difference regressions testing the impact of an AI technology shock on mutual funds' holdings. The treatment group consists of funds whose AI ratio was above the median in June 2016. The cutoff point is June 2017, when the transformer model was published. I choose three stock level measures and aggregate to fund level by calculating the weighted average. The three measures are: number of information event, market capitalization and age of stock. The independent variable is the AI ratio. All the independent variable and control variables are lagged for one quarter. A full description of all the variables can be found in the appendix. Standard errors are clustered at the fund family and quarter level; t-statistics are reported in parentheses.

	Information Event		Marketcap		Age of Stock	
	(1)	(2)	(3)	(4)	(5)	(6)
Post×Treatment	1.017***	0.872***	7.308***	6.564***	0.031	0.037
	(3.17)	(3.44)	(2.98)	(3.05)	(0.27)	(0.36)
Qret (%)		0.518***		1.853		0.001
		(4.18)		(1.88)		(0.09)
Logsize		0.551		12.314		0.235
		(0.62)		(1.47)		(1.39)
Logfamilysize		0.322		-0.315		0.229
		(0.63)		(-0.10)		(0.93)
Logage		-7.160***		-38.463***		-1.989***
		(-6.30)		(-3.89)		(-3.29)
Flow		-0.449		4.625		0.316
		(-0.73)		(0.95)		(0.76)
Expenses (%)		-2.065		-15.90**		0.588*
		(-1.31)		(-2.19)		(1.95)
Observations	5,124	5,000	5,123	4,999	5,071	4,947
R-squared	0.898	0.905	0.931	0.933	0.974	0.974
Fund FE	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES

Table 9: **Future Performance of Buy versus Sell in each AI Quintile**

This table presents the DGTW characteristic-adjusted performance of stocks that were either purchased or sold by mutual funds, which are sorted based on their AI ratio as described in Section 3.2. Mutual funds are sorted into 5 portfolios based on their AI measure at the beginning of each semi-year. I aggregate the holdings of each portfolio into a single portfolio. Within each portfolio, fund trades are further categorized into buy and sell trades, as outlined in Section 4.3. Additionally, all stocks are classified into two categorieshigh or low information intensitybased on whether the number of information events each quarter is above or below the median. Columns (1) to (3) report the time-series mean quarterly DGTW benchmark-adjusted returns of trades for stocks with information events above the median. Columns (4) to (6) report the time-series mean quarterly DGTW benchmark-adjusted returns of trades for stocks with information events below the median.

Quintile	Stock with more Infor Events			Stock with less Infor Events		
	(1)	(2)	(3)	(4)	(5)	(6)
	Buy	Sell	Difference	Buy	Sell	Difference
All Funds	0.0024	-0.0022*	0.0047**	0.0019	0.0022	-0.0002
	(1.53)	(-1.70)	(2.56)	(0.66)	(0.69)	(-0.21)
1(Low)	0.0008	0.0001	0.0006	0.0003	0.0042	-0.0039
	(0.39)	(0.11)	(0.27)	(0.09)	(1.16)	(-1.25)
2	0.0013	-0.0015	0.0028	0.0045	-0.0014	0.0059
	(0.82)	(-0.63)	(0.98)	(1.60)	(-0.38)	(1.60)
3	-0.0001	-0.0003	0.0002	-0.0007	-0.0008	0.0002
	(-0.04)	(-0.15)	(0.07)	(-0.20)	(-0.24)	(0.06)
4	0.0012	-0.0026	0.0038*	0.0040	0.0021	0.0019
	(0.66)	(-1.63)	(1.72)	(1.36)	(0.58)	(0.59)
5(High)	0.0062**	-0.0049*	0.0111**	0.0014	0.0036	-0.0023
	(2.21)	(-1.74)	(2.49)	(0.42)	(1.09)	(-0.91)

Table 10: **AI and Manager Skill**

This table reports the regression results of the manager skills (Equation 15 in Section 4.3). The dependent variables are *Timing* and *Picking*, defined in equations (13) and (14). The independent variable is the AI ratio. All the independent variable and control variables are lagged for one quarter. Standard errors are clustered at the fund family and quarter level; t-statistics are reported in parentheses.

	Stock Picking		Market Timing	
	(1)	(2)	(3)	(4)
AI ratio (%)	0.028***	0.075**	0.012	0.019
	(3.07)	(2.03)	(0.51)	(0.97)
Qret(%)		-0.065		0.017
		(-1.23)		(0.51)
Logsize		-0.546***		-0.113
		(-3.10)		(-1.41)
Logfamilysize		-0.130		0.046**
		(-1.17)		(2.32)
Logage		0.180		0.135
		(0.68)		(1.19)
Flow		-0.362		-0.173
		(-1.57)		(-0.94)
Expense		-0.235		0.001
		(-1.30)		(0.00)
Observations	35,999	34,889	36,031	34,921
R-squared	0.154	0.165	0.905	0.903
Fund FE	YES	YES	YES	YES
Time FE	YES	YES	YES	YES

Table 11: **Manager Turnover**

This table reports the impact of AI on mutual fund manager turnover. The dependent variable is the manager turnover ratio in Column (1) to (2) and the manager turnover dummy in Column (3) to (6) The independent variable is the AI ratio. I use probit regression in Column (5) and Column (6). Standard errors are clustered at the fund family and quarter level; t-statistics are reported in parentheses.

Dependent Variable	Manager Turnover		Manager Turnover Dummy			
	(1)	(2)	(3)	(4)	(5)	(6)
AI ratio (%)	0.000	0.001	-0.000	0.001	-0.003	0.018
	(0.24)	(0.85)	(-0.15)	(0.81)	(-0.12)	(0.77)
Pastyearalpha (%)		-0.000		-0.000		-0.000
		(-1.08)		(-0.82)		(-0.62)
Logsize		-0.003***		-0.004***		-0.044***
		(-5.07)		(-3.50)		(-4.10)
Logfamilysize		0.002***		0.004***		0.027**
		(5.17)		(3.38)		(2.35)
Logage		0.002		-0.002		-0.027
		(1.18)		(-0.52)		(-0.80)
Team Size		0.002***		0.017***		0.115***
		(3.49)		(8.45)		(18.27)
Observations	35,962	30,976	35,962	30,976	35,962	30,976
R-squared	0.004	0.007	0.006	0.033		
Time FE	YES	YES	YES	YES	YES	YES

A Variable Definitions

Name	Definition	Sources
AI ratio	The measure of AI intensity of a fund (fund company), calculated by the AI labor stock divided by the total labor stock	Burning Glass
Flow	Net flow into the fund/share class divided by lagged T-NA, adjusted by M&A (Equation 1)	CRSP
Activeshare	The active share calculated from the fund holding following Doshi et al. (2015).	Refinitiv
Alpha	Return of fund adjusted by Carhart four factor model	CRSP
Pastyearalpha	The alpha of a fund in the past one year	CRSP
Logage	The natural logarithm of fund age (in year)	CRSP
Logsize	The natural logarithm of fund TNA	CRSP
Logfamilysize	The natural logarithm of fund familu total TNA	CRSP
Turnover	The turnover rate of a fund/share class	CRSP
Expenses	The expense ratio of a fund/share class	CRSP
Holding Marketcap	The weighted average marketcap of the holding of a fund	Refinitiv&CRSP
Holding Information Event	The weighted average number of information events of the holding of a fund	Capital IQ
Holding Age of Stock	The weighted average firm age of the holding of a fund	Refinitiv&CRSP
Stocking Picking	The stock picking skill of a manager (Equation 13)	Refinitiv&CRSP
Market Timing	The market timing skill of a manager (Equation 14)	Refinitiv&CRSP
Team Size	The number of mutual fund managers in a fund.	Morningstar
Manager turnover	A variable equals to 1/Team Size if a fund manager leaves in that quarter.	Morningstar
Manager turnover dummy	Dummy that euqals one if the a fund manager leaves in that quarter.	Morningstar

B Indentify AI jobs

Table 12: **Examples of skills with high AI score and low AI score**

This table shows some examples of skills from Burning Glass job postings. The two leftmost columns display 20 skills with high AI scores, while the two rightmost columns display 20 skills with low AI scores. The AI scores of the corresponding skills are also reported.

High AI skill	AI score	Low AI skill	AI score
Artificial Intelligence	1	Credit Risk	0.01917
Machine Learning	1	Risk Management	0.018852
Natural Language Processing	1	Workflow Management	0.018774
Data Science	0.494301	Change Management	0.018208
Unstructured Data	0.466061	Regulatory Compliance	0.018174
Scala (Programming Language)	0.32161	Equities	0.017872
Algorithms	0.29445	Asset Management	0.01779
R (Programming Language)	0.285907	Decision Making	0.017678
Big Data	0.28414	Portfolio Management	0.01756
Data Engineering	0.268801	Microsoft Excel	0.017403
Advanced Analytics	0.214577	Risk Mitigation	0.017251
Statistical Modeling	0.210637	Risk Appetite	0.017209
Distributed Computing	0.188049	Management	0.017078
Apache Kafka	0.187259	Finance	0.016926
Python (Programming Language)	0.185729	Investments	0.016673
MATLAB	0.175258	Accountability	0.016358
Data Mining	0.1621	Project Management	0.016309
Applied Mathematics	0.157066	Internal Auditing	0.015974
Model Risk Management	0.15424	Leadership	0.015913
Statistics	0.151451	Sales Prospecting	0.013233

Table 13: **Examples of jobs with high AI score before cleaned by GPT**

This table shows ten examples of jobs with high AI scores in Burning Glass job postings. The four columns report the company name, skill requirements, job title, and the AI score. These jobs have not been evaluated by GPT-4 yet.

Company Name	Job Skills	Job Title	AI Score
Truist Financial	Machine Learning	ML Default Support Specialist II	1.00
JPMorgan Chase	Machine Learning	Consumer & Community Banking - Card Risk Machine Learning - Sr. Associate	1.00
Bank of America	Artificial Neural Networks Unsupervised Learning Machine Learning Algorithms Machine Learning TensorFlow Deep Learning Artificial Intelligence Data Analysis	Data Scientist - Machine Learning/AI/Python	0.68
Fidelity Investments	Algorithms Python (Programming Language) Knowledge Graph Research Papers Reinforcement Learning Data Analysis Question Answering Machine Learning Chatbot TensorFlow Conversational AI Deep Learning Natural Language Processing Apache MXNet Elasticsearch Keras (Neural Network Library) Artificial Intelligence	Senior Data Scientist	0.55
BlackRock	Equities Python (Programming Language) Portfolio Optimization Mathematical Modeling Mathematics Statistics Machine Learning Natural Language Processing	Vice President, Systematic Active Equity Team	0.31

Continued on next page

Company Name	Job Skills	Job Title	AI Score
Goldman Sachs	Probability And Statistics Financial Modeling Machine Learning Algorithms Statistics Program Process Monitoring Portfolio Management Machine Learning Java (Programming Language) Predictive Modeling TensorFlow Natural Language Processing Production Process Scripting PyTorch (Machine Learning Library) Keras (Neural Network Library) Computational Statistics Risk Modeling Economics Computer Science Credit Risk Modeling Mathematical Finance R (Programming Language)	Transaction Banking Data Scientist / Quantitative Engineer Lending Associate	0.29
BlackRock	Portfolio Management Python (Programming Language) Financial Economics Mathematics RStudio Econometrics Natural Language Processing MATLAB Economics Artificial Neural Networks	Associate, Portfolio Manager	0.25
The Vanguard Group	Finance Fixed Income Equities Python (Programming Language) Machine Learning Amazon Web Services Artificial Intelligence Advanced Analytics Investment Management Risk Management	ML Engineer - Investment (Python/AWS)	0.24

Continued on next page

Company Name	Job Skills	Job Title	AI Score
Morgan Stanley	Asset Allocation Asset Classes Research Presentations Procurement Statistical Software Environmental Social And Corporate Governance (ESG) Artificial Intelligence Forecasting Statistical Programming Data Strategy R (Programming Language)	ESG Data Specialist	0.13
T. Rowe Price Group	Object-Oriented Programming (OOP) Research Management Application Programming Interface (API) Microservices Application Development Automation Unix Tooling Consensus Protocol Linux Mentorship Multi-Tenant Cloud Environments Systems Development Life Cycle Python (Programming Language) Prometheus (Software) Java (Programming Language) Hybrid Cloud Computing Operations Observability Scalability Scripting Amazon Web Services Business Strategies Artifactory Site Reliability Engineering Grafana	Lead Site Reliability Engineer	0.08

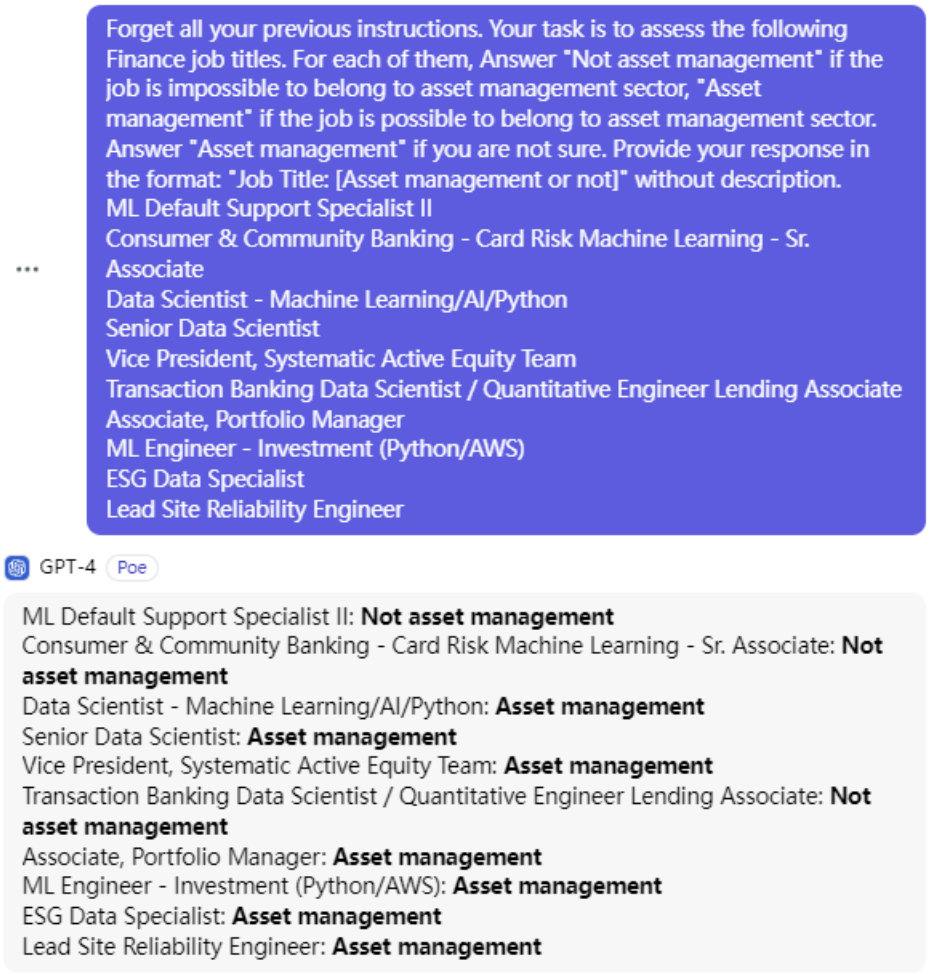


Figure 5: **Determine the Asset Management Industry** This figure illustrates how to determine whether a job belongs to the asset management sector or not using GPT-4. The content in blue represents the input, which includes ten jobs and the prompt. The content in white is the response from GPT-4. These ten jobs are the same as the ten examples in Table 10.

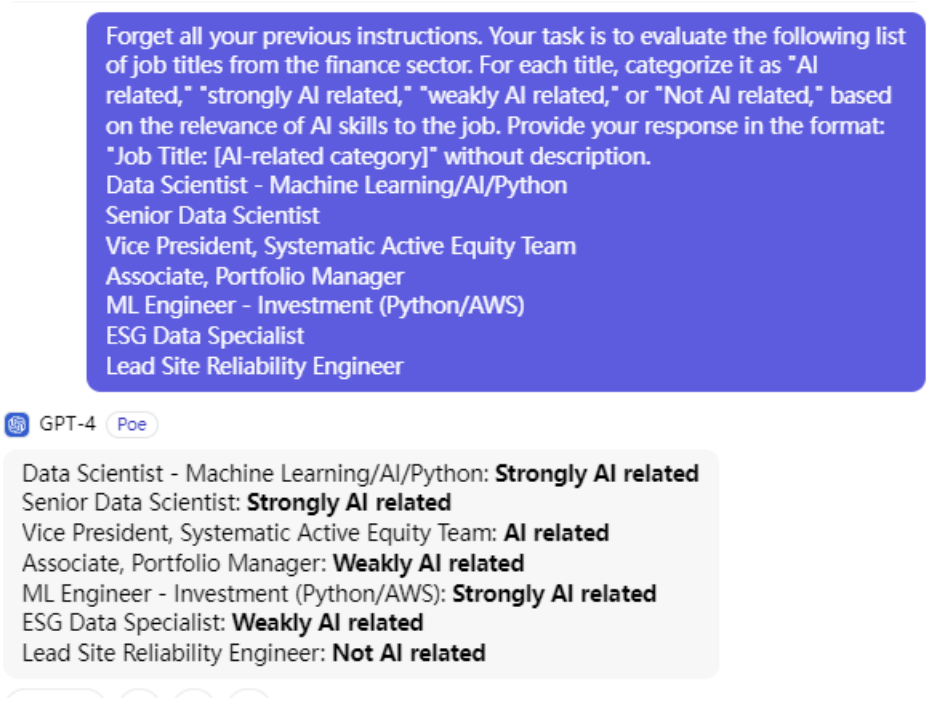
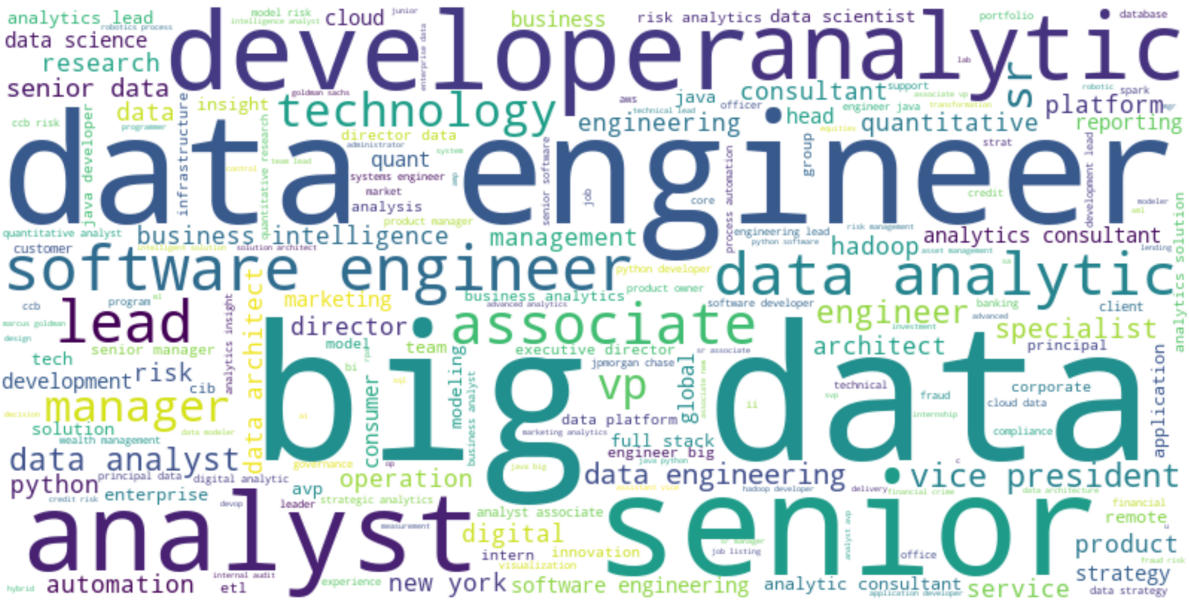


Figure 6: **Determine the AI Job** This figure illustrates how to determine whether a job is an AI job or not using GPT-4. The content in blue represents the input, which includes seven jobs and the prompt. The content in white is the response from GPT-4. These seven jobs are the same as the job in asset management industry in the previous figure.

Panel A: Strongly AI related



Panel B: Not AI related



Figure 7: **Keywords in the job titles of the categorized strongly AI related and not AI related** This figure shows the keywords in the job titles of different types of jobs categorized by GPT-4. Larger fonts indicate a higher word frequency.

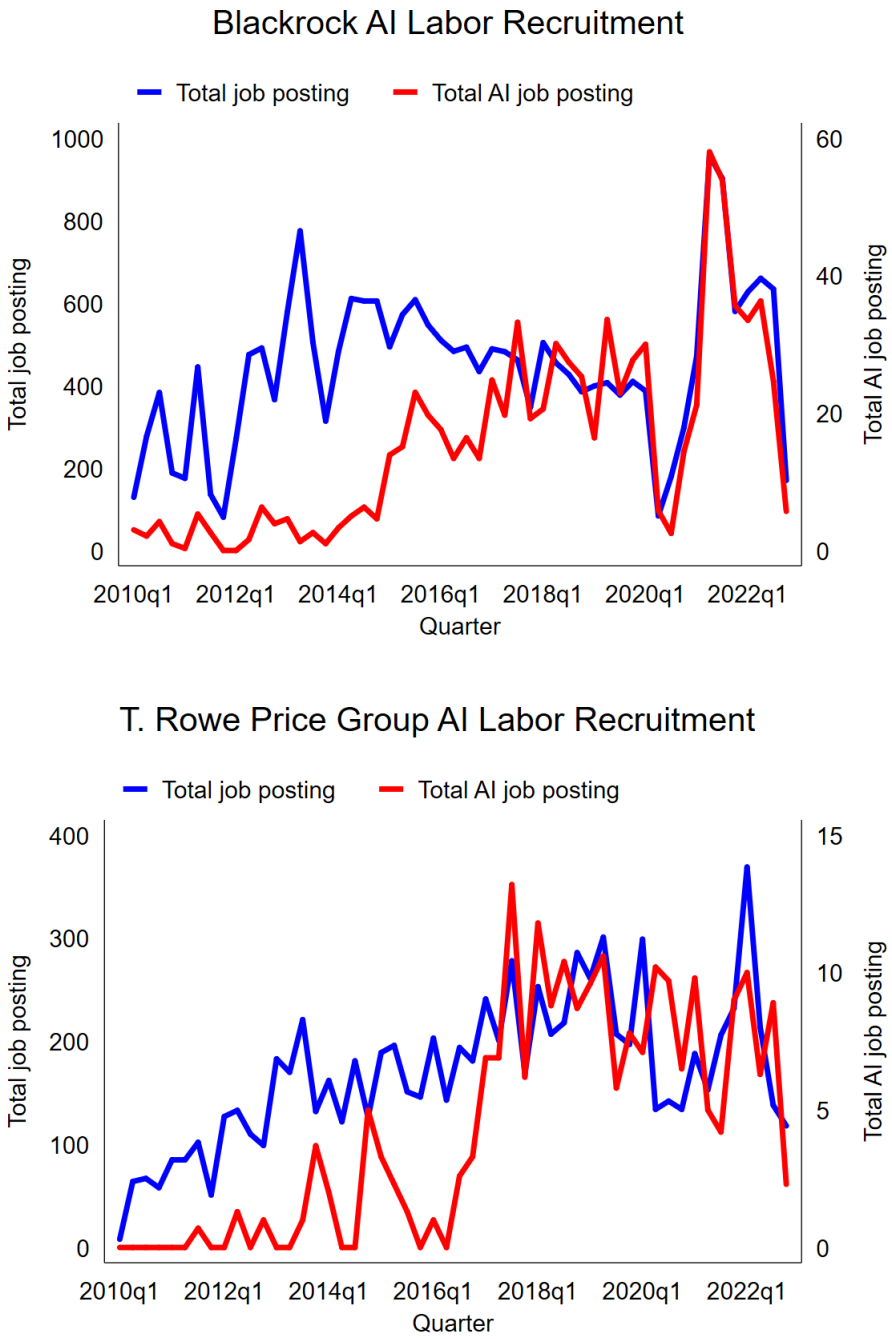


Figure 8: **AI Labor Recruitment: Examples** These figure plot the AI labor recruitment for two asset management companies (Blackrock and T. Rowe Price Group) quarterly. The red line is the total number of AI job posting for the fund company, correspond to the right y-axis. The blue line is the total number of all the job posting for that company, correspond to the left y-axis.

C Robustness Check: AI Measure

Table 15: **AI Measure Correlation**

This table reports the correlation matrix of AI measures using different cutoffs. The five AI measures are calculated with cutoffs equal to 0.07, 0.075, 0.08, 0.085, and 0.09, respectively. For example, a cutoff equal to 0.07 means that a job will be classified as a non-AI job if its AI score is less than 0.07.

	AI_0.07	AI_0.075	AI_0.08	AI_0.085	AI_0.09
AI_0.07	1				
AI_0.075	0.999	1			
AI_0.08	0.996	0.999	1		
AI_0.085	0.992	0.996	0.998	1	
AI_0.09	0.990	0.994	0.996	0.999	1

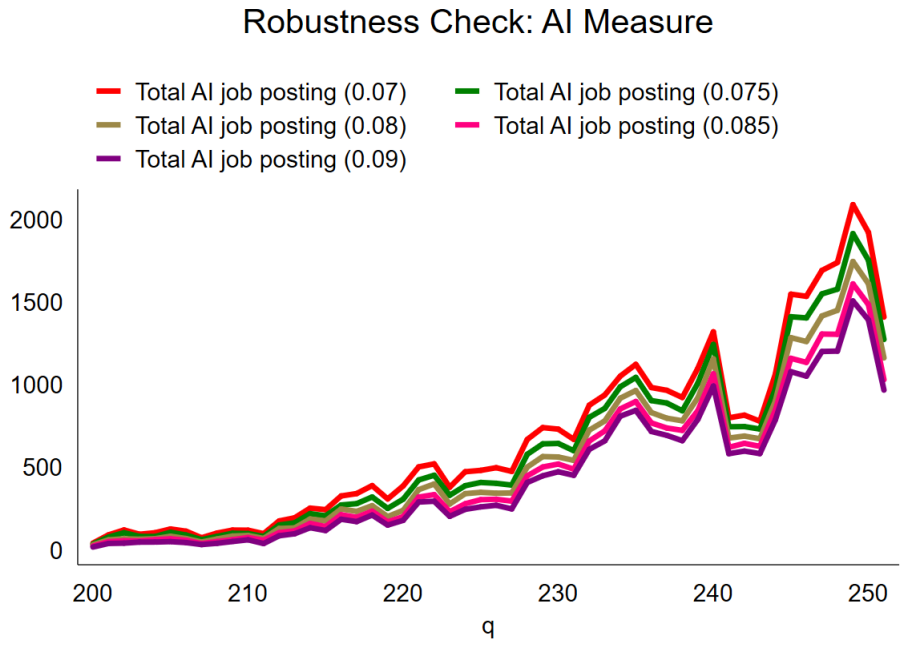


Figure 9: **Robustness Check: AI Measure** This figure plots the total number of AI job posting using different cutoffs. The five AI measures are calculated with cutoffs equal to 0.07, 0.075, 0.08, 0.085, and 0.09, respectively. For example, a cutoff equal to 0.07 means that a job will be classified as a non-AI job if its AI score is less than 0.07.

D AI Measure and Fund Characteristic

Table 16: **The relationship between AI and other variables**

This table reports the results of regressing different fund variables on lagged AI ratio. The dependent varibales are flow, fee, fund age, turnover and Activeshare, respectively. Standard errors are clustered at the fund family and quarter level; t-statistics are reported in parentheses.

	(1)	(2)	(3)	(4)	(5)
Dependent Variables	Flow	Expenses	Fund Age	Turnover	Activeshare
AI ratio (%)	0.002*	-0.000	-0.172	0.010	-0.007
	(1.79)	(-0.65)	(-0.45)	(0.31)	(-1.20)
Observations	39,509	39,698	39,698	29,389	31,079
R-squared	0.009	0.074	0.014	0.019	0.004
Time FE	YES	YES	YES	YES	YES