Generative AI and Asset Management*

Jinfei Sheng† Zheng Sun‡ Baozhong Yang§ Alan Zhang¶

March 2025

Abstract

This paper proposes a novel measure of investment companies' reliance on generative AI, focusing on its implications for hedge funds. We document a sharp increase in generative AI usage by hedge funds after ChatGPT's 2022 launch. A difference-in-differences test shows that hedge funds adopting generative AI earn 3-5% higher annualized abnormal returns than non-adopters. We further identify this effect by exploiting ChatGPT outages as exogenous shocks. The outperformance originates from funds' AI talent and ChatGPT's strength in analyzing firm-specific information. Non-hedge funds yield no significant returns from AI adoption, suggesting generative AI may widen existing disparities among investors.

JEL Classification: C81, G11, G14, G23

Keywords: Generative AI, ChatGPT, Hedge Funds, Reliance on AI information (RAI), Portfolio Return, Alpha, Outage, AI Disparity

*We are grateful to Vikas Agarwal, Maxime Bonelli, Hui Chen, Yong Chen, Winston Dou, Andrea Eisfeldt, Thierry Foucault, Itay Goldstein, Jiekun Huang, Shiyang Huang, Tingliang Huang (Discussant), Hee-Seo Han, TC Lin, Fangzhou Lu, Ben Lourie, Stijn van Nieuwerburgh, Shumiao Ouyang (Discussant), Devin Shanthikumar, Qiguang Wang, Wanyi Wang, Franco Wong, Ayako Yasuda for helpful discussions and suggestions, and to conference and seminar participants at 2024 SFS Cavalcade Asia-Pacific Conference, 2024 ABFER-JFDS Conference on AI and FinTech, Australian Finance and Banking Conference, 2024 Winter Institute of AISSEAI, CEIBS, DUFE, Georgia State University, Iowa State University, Korea University, Shanghai Jiaotong University, University of California-Davis, University of California-Irvine, University of Hong Kong, University of Miami, and York University for helpful comments and suggestions.

†Merage School of Business, University of California, Irvine, Email: jinfei.sheng@uci.edu

‡Merage School of Business, University of California, Irvine, Email: zhengs@uci.edu

§J. Mack Robinson College of Business, Georgia State University, Email: bzyang@gsu.edu

¶College of Business, Florida International University, Email: alan.zhang@fiu.edu

Generative AI and Asset Management

March 2025

Abstract

This paper proposes a novel measure of investment companies' reliance on generative AI, focusing on its implications for hedge funds. We document a sharp increase in generative AI usage by hedge funds after ChatGPT's 2022 launch. A difference-in-differences test shows that hedge funds adopting generative AI earn 3-5% higher annualized abnormal returns than non-adopters. We further identify this effect by exploiting ChatGPT outages as exogenous shocks. The outperformance originates from funds' AI talent and ChatGPT's strength in analyzing firm-specific information. Non-hedge funds yield no significant returns from AI adoption, suggesting generative AI may widen existing disparities among investors.

JEL Classification: C81, G11, G14, G23
Keywords: Generative AI, ChatGPT, Hedge Funds, Reliance on AI information (RAI), Portfolio Return, Alpha, Outage, AI Disparity

1. Introduction

In the asset management industry, information is the key to success. According to Grossman and Stiglitz (1980), sophisticated investors earn alphas by engaging in costly searches for new information and by accurately processing it in a timely manner. However, effectively doing so is challenging due to the vast amount and complexity of potentially useful information for asset pricing (Chen, Cohen, Gurun, Lou, and Malloy, 2020; Martin and Nagel, 2022). Artificial intelligence (AI) has experienced substantial advancement in the past two decades, leading to vast adoptions of the technology by companies to process data and aid their decision-making.[1] However, AI has been highly technical and its applications require special talents, which leads to a scarcity of human capital in this area[2] and a challenge in generating returns on investment with AI.[3]

Generative AI, exemplified by ChatGPT, is a significant, disruptive revolution in AI techniques. Their performance in understanding texts, solving problems, and producing answers is truly remarkable and comparable to or exceeds human performance.[4] More importantly, different from previous AI tools, generative AI does not require complicated training and tuning and can be intuitively used by the general public, leading to their rapid adoption, e.g., ChatGPT is the fastest app to reach 100 million users.[5] Given the potential of generative AI, understanding how it is used by investors and its impact on investing thus can have important implications. However, such studies are challenging due to the lack of observable data on the use of generative AI by companies and investors.

[1]See, for example, Webb (2019), Acemoglu, Autor, Hazell, and Restrepo (2022), Babina, Fedyk, He, and Hodson (2024), and Abis and Veldkamp (2024).

[2]Sources: "AI talent war on Wall Street hits Goldman Sachs hardest," November 28, 2023, William Shaw, *Fortune*; "AI talent war heats up in Europe," March 11, 2024, Martin Coulter, *Reuters*; "Inside Silicon Valley's AI talent war," March 28, 2024, *Wall Street Journal* Podcast

[3]Sources: "Can an A.I. hedge fund beat the market?" August 25, 2020, Jeremy Kahn, *Fortune*; "Hedge funds find it's really hard to beat the market With AI," October 6, 2023, Justina Lee, *Bloomberg*.

[4]Sources: "ChatGPT passes exams from law and business schools," January 26, 2023, Samantha Murphy Kelly, *CNN*; "M.B.A. students vs. AI: Who comes up with more innovative ideas?" September 9, 2023, Christian Terwiesch and Karl Ulrich, *Wall Street Journal*.

[5]Source: "ChatGPT sets record for fastest-growing user base," February 2, 2023, Krystal Hu, *Reuters*.

In this paper, we propose a novel approach to measure the reliance on generative AI of investment companies and apply this measure to study the impact of generative AI on the asset management industry. In our study, we focus mostly on hedge funds since they are typically regarded as the most informative investors and earliest adopters of new technologies.[6] We propose to address the following research questions: Are generative AI technologies widely adopted by hedge fund companies? Does such adoption affect their performance? Does the availability of this intuitive tool help to level the playing field of the asset management industry?

To construct our measure of generative AI adoption, *Reliance on AI Information* (*RAI*), we utilize the 13F quarterly trades of hedge fund companies. We consider two types of information that correlate with trades of hedge fund companies: financial variables about firm fundamentals and information generated by ChatGPT based on conference calls (i.e., AI information). *RAI* measures that given the existing financial variables, what additional percentage of the variation in fund portfolio composition can be explained by AI information. In other words, *RAI* captures the degree to which fund managers' portfolio decisions are influenced by AI-generated information in addition to the existing set of fundamental variables.

Specifically, we follow a two-step procedure as in Kacperczyk and Seru (2007). In the first step, we look into the explanatory power (i.e. R-squared) of financial variables on hedge fund companies' trades. Next, we calculate the incremental explanatory power when adding AI-generated information. *RAI* is estimated as the incremental R-squared through this procedure. The measure is closely related to the coefficient of partial determination, which is commonly used to measure the marginal contribution of new variables when other variables have been included in the model.

Our *RAI* measure has two advantages. First, by capturing the marginal contribution

[6]For example, a 2018 BarclayHedge survey of hedge fund managers finds that more than half of hedge funds use AI and machine learning in their investment strategies. Source: "Majority of hedge fund pros use AI/machine learning in investment strategies," July 17, 2018, *BarclayHedge*.

of AI information to hedge funds' portfolio change, the measure identifies the usage by portfolio managers for investment analysis purposes, rather than other reasons. Second, our methodology can be applied to all investment companies with holdings information, allowing us to conduct a systematic analysis of the effect of generative AI on their performance.

Using the *RAI* measure, we first examine the adoption of generative AI among hedge funds. The time trends in *RAI* of hedge funds reveal a sharp and abrupt increase starting in 2022, coinciding with the introduction of the underlying base model of ChatGPT. This suggests substantial adoption of generative AI by hedge fund companies. To formally test the adoption of AI by hedge funds, we conduct a partial F-test, widely used in the literature (e.g., Greene, 2002, p101). We also estimate the false positive rate during this process.[7] Our tests indicate that the false positive rate is low and only around 2%. After adjusting for the false positive rate, we find that 19% of hedge funds started to adopt generative AI in 2022 at the significance level of 1%. This percentage remains around 18% in 2023. This notable adoption rate is consistent with the speed at which the general public embraces ChatGPT.[8]

We then investigate the characteristics of early adopters of ChatGPT. Ex ante, it is not clear what type of funds are more likely to use ChatGPT first. On the one hand, small hedge funds have incentives to use new tools to establish themselves and get an edge. Therefore, the AI tool may help level the playing field of the asset management industry. On the other hand, large funds have more resources to utilize new tools quickly. Our finding shows that large funds tend to adopt ChatGPT first. We also find that more active funds and funds with better past performance have a greater likelihood of becoming early

[7] One potential concern with our measure is that if a fund happens to obtain information that correlates with ChatGPT signals but does not actually use generative AI, then their *RAI* may be overestimated, generating false positives. We formally estimate the false positive rate using a partial F-test based on the pre-2022 sample period. The idea is that if a fund during this early period is estimated to have a significant *RAI*, it is likely to be a false positive case because the ChatGPT was not available to the general public then.

[8] A February 2024 Pew Research survey shows that 23% of American adults had utilized ChatGPT. Source: "Americans' use of ChatGPT is ticking up, but few trust its election information," March 26, 2024, Pew Research Center.

adopters.

We next study the relation between *RAI* and hedge fund performance. Results from panel regressions indicate that hedge funds with higher reliance on ChatGPT earn better returns during the post-ChatGPT period. We also conduct a difference-in-differences (DiD) test to examine whether *RAI*'s predictive power for fund performance significantly increases following the introduction of ChatGPT. Our tests show that hedge fund companies with a higher *RAI* generate significantly higher raw and risk-adjusted returns. The economic magnitude of this effect is large. An interquartile increase in the reliance on generative AI is associated with a gain of 3 to 5% in annualized abnormal returns, for different asset pricing models. Therefore, generative AI does bring substantial benefits to hedge fund companies that adopt this new technology.

To further identify the relation between the usage of generative AI and fund performance, we exploit ChatGPT outages as exogenous shocks to the availability of ChatGPT to its users. In a triple DiD test, we find the effect of *RAI* on returns is significantly more muted when ChatGPT services experience more outages, suggesting a causal relationship between AI adoption and hedge fund performance.

To the extent that generative AI is accessible to all, a natural question is whether it benefits all institutions in the asset management industry equally. We find this not to be the case. We examine non-hedge fund companies and find that their AI adoption does not result in significantly better returns. Furthermore, large and more active hedge fund companies are able to leverage generative AI to obtain significant returns, while small and passive firms fail to do so. Taken together, the evidence suggests that applying the intuitive AI tool productively still requires additional resources such as data and expertise. This also implies that generative AI could enlarge the disparity among investors rather than level the playing field.

Finally, we investigate the potential mechanisms of how generative AI helps with

asset management and conduct two sets of tests. First, we hypothesize the availability of human capital with AI expertise within the hedge fund company increases the effectiveness of adoption. Indeed, we find that hedge funds with AI-skilled human capital generate much higher performance from their use of generative AI, consistent with the notion that these AI talents can use the tools more effectively. Second, we examine whether generative AI can help funds better analyze certain information. For this purpose, we decompose the *RAI* measure into three components regarding macroeconomic, firm-policy, and firm-performance information. We find that only firm-level policy and performance information contributes to greater fund performance, indicating that generative AI is mostly useful for funds to analyze firm-specific information.

This paper contributes to several streams of literature. First, our research contributes to the literature on the skill and performance of hedge funds and mutual funds. Several studies document evidence of hedge fund and mutual fund skill through examination of their stock holdings, e.g., Wermers (2000); Kacperczyk, Sialm, and Zheng (2005, 2008); Griffin and Xu (2009); Agarwal, Jiang, Tang, and Yang (2013); Aragon, Hertzel, and Shi (2013). Furthermore, a number of studies identify characteristics that distinguish skilled hedge funds, such as strategy distinctiveness (Sun, Wang, and Zheng, 2012), risk exposure to systematic factors (Titman and Tiu, 2011), market timing (Chen and Liang, 2007), market liquidity timing (Cao, Chen, Liang, and Lo, 2013), exposure to investor sentiment (Chen, Han, and Pan, 2021), geographical preference (Sialm, Sun, and Zheng, 2020), and unobserved performance (Agarwal, Ruenzi, and Weigert, 2023). Different from these studies, our paper shows that the adoption of disruptive generative AI technology can also contribute substantially to fund performance.

Second, our paper is also related to the use and implications of new technologies and data in asset management, e.g., alternative data (Bonelli and Foucault, 2023), and AI in venture capital investment (Lyonnet and Stern, 2022; Bonelli, 2023). We complement this

literature by being the first to study the adoption of ChatGPT, a significant and disruptive revolution in AI technologies, in the asset management industry. Our construction of a unique generative AI reliance measure enables the study of the implications of this disruptive AI technology.

Third and more generally, our paper contributes to the literature on the impact of AI on the economy and financial markets.[9] Theoretically, AI may come with unexpected costs such as lower price efficiency (Dugast and Foucault, 2023; Colliard, Foucault, and Lovo, 2022; Dou, Goldstein, and Ji, 2023). Empirically, there is evidence on both the positive and negative sides of AI. For instance, Fuster, Goldsmith-Pinkham, Ramadorai, and Walther (2022) show the negative effects of AI on the disparity in the credit markets. Cao, Jiang, Wang, and Yang (2024) find that human wisdom and AI power complement each other in stock analyses.[10]

Different from previous AI advances such as machine learning, generative AI represents a major, unexpected breakthrough in AI technologies that first makes AI widely available to the public and investors with low costs. Most related to our paper are several very recent studies that examine the effects of language language models (LLM)/generative AI on stock prices and job markets (Eisfeldt, Schubert, and Zhang, 2023), and corporate customer service quality (Brynjolfsson, Li, and Raymond, 2023). Such studies are challenging to conduct in general due to the difficulty in obtaining data on the use of generative AI by companies. We contribute to the literature by conducting the *first large-scale study* of the use of generative AI in the asset management industry. The setting of investment companies, the availability of holdings data, and our methodology allow us to infer the use of generative AI and study its implications. Our findings reveal that despite its accessibility,

[9] One important early application of AI in the finance industry is robo-advising, which can improve retail investors' welfare (D'Acunto, Prabhala, and Rossi, 2019; Rossi and Utkus, 2024). Algorithmic aversion, however, can hinder AI adoption (Greig, Ramadorai, Rossi, Utkus, and Walther, 2022).

[10] Relatedly, AI affects the real economy such as workforce composition (Babina, Fedyk, He, and Hodson, 2023). Also, data management affects the workforce in the financial services industry (Abis and Veldkamp, 2024).

generative AI may in fact further increase disparities among market participants. This carries implications as society prepares to widely adopt generative AI technologies.[11]

Finally, there is an emerging literature that applies generative AI techniques to research in finance and economics, e.g., evaluating news sentiment (Lopez-Lira and Tang, 2023), classifying Federal Reserve policy stances (Hansen and Kazinnik, 2023), identifying lengthy discussions in earnings transcripts (Kim, Muhn, and Nikolaev, 2023), quantifying information content in answers (Bai, Boyson, Cao, Liu, and Wan, 2023), understanding expected corporate policies and the implications on asset prices (Jha, Qian, Weber, and Yang, 2023), and analyzing corporate culture and its impact (Li, Mai, Shen, Yang, and Zhang, 2023).[12] This literature utilizes the power of generative AI to perform deep analysis of textual data and expand the horizon of economic research.[13] While we also rely on the use of generative AI in the definition of our key measure, we have a distinct focus on studying the implications of AI adoption in the asset management industry.[14]

2. Institutional Background and Data

In this section, we describe the institutional background of the history and development of ChatGPT, as well as the datasets we use.

2.1. Background on ChatGPT

Developed by OpenAI, ChatGPT (Chat Generative Pre-trained Transformer) represents a significant milestone in natural language processing and AI. The underlying technology of ChatGPT is based on the Transformer architecture of deep learning models (Vaswani

[11] "Business Schools Are Going All In on AI," April 3, 2024, Lindsay Ellis, *Wall Street Journal*.

[12] See also Korinek (2023) for a discussion of use cases of generative AI in economics research.

[13] A closely related branch of literature has applied large-language models and their foundation – the transformer models – in economic research, e.g., Cong, Tang, Wang, and Zhang (2021), Acikalin, Caskurlu, Hoberg, and Phillips (2022), Jiang, Kelly, and Xiu (2022).

[14] More generally, our paper is also related to the literature on textual analysis in finance (e.g., Tetlock, 2007; Loughran and McDonald, 2011; Hoberg and Phillips, 2010, 2016; Fisher, Martineau, and Sheng, 2022; Garcia, Hu, and Rohrer, 2023).

et al., 2017), which allows self-attention mechanisms, self-supervised training, and superior performance. Since 2018, Open AI has released increasingly capable Transformer-based pre-trained models, including GPT 1 in 2018, GPT 2 in 2019, and GPT 3 in 2020. GPT 3 was able to complete writing tasks in a much more polished manner than prior versions. These models serve as the predecessors of ChatGPT. OpenAI also made its Application Programming Interface (API) publically available in 2021.[15]

ChatGPT is based on the GPT 3.5 model series, which significantly increases the capabilities of prompt understanding and question answering. The first model in the GPT 3.5 series was released in March 2022 and became publicly available through the API platform.[16] The GPT 3.5 model was then further fine-tuned to produce ChatGPT 3.5 and formally launched to the public through a chat-based interface on November 30, 2022.[17]

ChatGPT is built upon a robust foundation of deep learning and AI advancements. The evolution from GPT-3 to ChatGPT 3.5 involved enhancements in model architecture, training data, and fine-tuning methodologies, including reinforcement learning with human feedback (RLHF). With increased parameters and improved algorithms, ChatGPT 3.5 exhibits far superior performance in understanding and generating human-like text responses across diverse contexts relative to earlier models. Furthermore, ChatGPT exhibits "emergent abilities" that allow it to tackle even problems in unfamiliar domains. As a result, ChatGPT took the world by surprise and made a remarkable debut, swiftly gaining popularity. By December 4, 2022, ChatGPT had over one million users. Subsequently, in January 2023, it reached a milestone of over 100 million users, positioning it as one of the fastest-growing consumer applications to date.[18]

After the initial release, Open AI made continual improvements to ChatGPT. For

[15]Source: "OpenAI's API now available with no waitlist," November 18, 2021, OpenAI.
[16]Source: "New GPT-3 capabilities: Edit & insert," March 15, 2022, OpenAI.
[17]See https://platform.openai.com/.
[18]According to a February 2024 Pew Research poll, 23% of American adults had tried ChatGPT. Source: "Americans' use of ChatGPT is ticking up, but few trust its election information," March 26, 2024, Pew Research Center.

example, it released ChatGPT Plus on February 1, 2023, which allows subscribers to access the most recent models and features. On March 1, 2023, Open AI made ChatGPT available through its API services. The latest and most advanced version, ChatGPT 4 was released on March 14, 2023. Figure 1 shows the timeline of the development of ChatGPT.

[Insert Figure 1 Here]

Admittedly, there are other generative AI tools beyond ChatGPT, such as Claude 3 by Anthropic and Llama 3 by Meta. We focus on ChatGPT for at least two reasons. First, ChatGPT was the first powerful large language model tool, allowing us to have a relatively longer sample period. Second, ChatGPT is arguably the most widely used generative AI tool by the public, including professional investors. Thus, it is intuitive to use ChatGPT in this setting. Also, to the extent that other generative AI tools generate signals correlated with ChatGPT, our measure can be viewed as a proxy for hedge funds' use of generative AI tools in general.[19]

2.2. Data: AI-generated Signals

The data used in this study come from various sources. ChatGPT is utilized to generate AI-predicted information about public firms from conference call transcripts.[20] Specifically, ChatGPT is queried with questions about firms' future policies in various areas, such as investment, employment, etc. For instance, one question we ask is *"Over the next quarter, how does the firm anticipate a change in its employment."* ChatGPT will answer this question based on earnings conference call transcripts. The set of questions is based on those in Jha, Qian, Weber, and Yang (2023, 2024). There are a total of 14 signals, or *GPT Score*s, generated by ChatGPT, covering firms' expectations about macroeconomic, industry, and

[19]Some hedge funds may want to use their proprietary generative AI models rather than ChatGPT. However, industry reports suggest that it takes a long time to develop a high-quality generative AI model customized to the financial industry needs (Source: "Finding value in generative AI for financial services," MIT Technology Review, 2023). Therefore it is unlikely that such a model was immediately available during the first few months of ChatGPT, the sample period that our paper focuses on.

[20]We thank the coauthors of Jha, Qian, Weber, and Yang (2024) for sharing data with us.

firm-specific performance and policy outcomes. A full list of questions can be found in Table A2 in the Appendix.

We focus on information from earnings conference calls for two reasons. First, it is well-documented that this data source is important for investors and well-accepted in the finance literature (e.g., Li, Mai, Shen, and Yan, 2021; Li, Mai, Shen, Yang, and Zhang, 2023). Second, Jha, Qian, Weber, and Yang (2023) show that signals from ChatGPT are high-quality and can be used to predict firms' future corporate policies and returns. Thus, we hypothesize that hedge funds may use ChatGPT to analyze earnings conference call texts to help with their investment decisions. Also, adding other information sources may increase the magnitude of the effect of AI tools. Thus, our estimate from earnings conference calls provides a lower bound of the generative AI's impact.

It is important to point out that the AI-generated information based on call texts may not necessarily reflect new information, but AI can still help fund managers process a large amount of unstructured data with forward-looking predictions. For this study, it does not matter whether the information from ChatGPT is new or not, as such information can aid managers' investment decisions in either case.

2.3. Data: Other Variables

Other data sources include Institutional (13F) Holdings from Thomson Reuters/Refinitiv, fundamental and market information data about portfolio firms from CRSP, Compustat, and I/B/E/S, and manual classification of 13F investment companies that operate hedge funds.[21]

We calculate the portfolio returns in quarter $t+1$ for each investment company i, based on its 13F holdings at the end of quarter t. *Return* is defined as the weighted average cumulative monthly return across all holdings in quarter $t+1$, where the weight is the

[21]The classification is based on several sources, including online business name datasets such as Bloomberg, company websites, and Form ADVs filed by investment companies. Our classification method is based on Agarwal, Jiang, Tang, and Yang (2013) and extends to recent years.

value of stock j held by i at the end of quarter t divided by the total value of all stocks held by i at the end of quarter t. We also calculate weighted average risk-adjusted returns using CAPM, the Fama-French three-factor model, and the Fama-French-Carhart four-factor model (Fama and French, 1993; Carhart, 1997). Take *CAPM Alpha* as an instance, at the end of quarter t, we use the monthly stock returns in the past 36 months to estimate the beta on the risk factor and calculate abnormal return as the difference between realized stock return minus stock return estimated with beta. *CAPM Alpha* is the weighted average cumulative monthly abnormal return across all holdings. *FF3 Alpha* and *FF4 Alpha* are constructed analogously.

In addition, we control for investment companies and their holdings characteristics. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since an investment company's first 13F report. *Turnover* is the minimum of purchases and sales over average total holdings values of the current quarter and the previous quarter, following Carhart (1997). *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Past Return* is the portfolio return in the previous quarter.

3. Reliance on AI Information (*RAI*)

In this section, we discuss how we construct our measure of the reliance on AI information. We also discuss the pros and cons of this measure and time trends. We then provide a formal test of AI adoption by hedge funds. Finally, we show the characteristics of early adopters of this technology.

3.1. *RAI*: Measure Construction

To measure the reliance on AI by hedge funds, we estimate the responsiveness of a hedge fund manager's portfolio changes to AI-predicted signals. We call this measure *Reliance on AI information* (*RAI*). For AI-generated information, we obtain ChatGPT-predicted signals

as in Jha, Qian, Weber, and Yang (2023). We construct *RAI* using a methodology similar to Kacperczyk and Seru (2007), who measure a fund's reliance on public information. Specifically, we estimate *RAI* using a two-step procedure. In the first step, at the end of each quarter t and for each investment company i, we run the following two regression models across the investment company's holdings changes in quarter t:

$$HoldingChange_{i,j,t} = \gamma_{i,t} \cdot X_{j,t-1} + \varepsilon_{i,j,t} \tag{1}$$

$$HoldingChange_{i,j,t} = \Sigma_{j=1}^{J} \beta_{i,t} \cdot GPT\ Score_{j,t-1} + \gamma_{i,t} \cdot X_{j,t-1} + \varepsilon_{i,j,t} \tag{2}$$

where $HoldingChange_{i,j,t}$ denotes a percentage change in split-adjusted holdings of stock j held by an investment company i from time t - *1* to t.[22] $X_{j,t-1}$ is a host of financial variables about firm fundamentals at the end of quarter $t-1$, including market capitalization, book-to-market, return on assets, stock return, and change in the analyst recommendation consensus. Note that analyst recommendation is an aggregate outcome of analysts' research based on various information sources in the public domain, including the earnings conference call transcripts. This variable has been used by the literature to capture information in the public domain (eg., Kacperczyk and Seru, 2007).[23] Thus, we benchmark against the public information available to fund managers without the deployment of generative AI tools.

GPT Score includes 14 signals generated by ChatGPT, covering firms' expectations about macroeconomic, industry, and firm-specific performance and policy outcomes. Note that the cross-sectional regressions are conducted separately for each investment company i and quarter t. We define the R^2 from equation (1) as $R^2_{fundamental,i,t}$, and the R^2 from (2) as $R^2_{AI,i,t}$.

[22] Adding a new stock position would imply an infinite increase, so we set $HoldingChange_{i,j,t}$ to 100% for these cases, following Kacperczyk and Seru (2007).

[23] We use information at the end of quarter t-1 to ensure that such information is available to fund managers when they make portfolio changes during quarter t.

In the second step, we define *RAI* of investment company i at time t as the difference between these two R^2s, which is presented as follow:

$$RAI_{i,t} = R^2_{AI,i,t} - R^2_{fundamental,i,t} \tag{3}$$

The incremental R^2 estimated through this procedure is closely related to the coefficient of partial determination, which is commonly used to measure the marginal contribution of a new variable when other variables have been included in the model. Intuitively, *RAI* quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT in addition to the existing set of fundamental variables. As a robustness check, we also create an alternative reliance measure *Alt. RAI*, defined as the difference in R^2 (i.e., *RAI*) scaled by R^2 in equation (1) and report the results in section 6.2.

Given that generative AI is a recent phenomenon, we restrict our sample period from 2016Q1 to 2023Q2.[24] Our final sample consists of 633 unique hedge fund companies and 10,762 company-quarter observations. Table 1 reports the summary statistics for key variables of interest and control variables. Overall, our sample exhibits significant cross-sectional variation in *RAI*. The average value of *RAI* equals 0.26, with a standard deviation of 0.216 and an inter-quartile spread of 0.320.

[Insert Table 1 Here]

3.2. Pros and Cons of *RAI*

Ideally, researchers want to know whether a hedge fund actually uses ChatGPT to aid their portfolio decision, for each time period and each hedge fund in the market. However, such data is challenging to obtain for at least two reasons: First, suppose one can possibly observe the subscription data of ChatGPT web or API services, we still do not

[24] TR 13F holdings data are available until 2023Q2 as of the writing of this draft.

observe whether ChatGPT is used for investment analysis or some other purposes such as generating marketing materials or copy editing. Second, one can conduct a survey to ask what a hedge fund uses ChatGPT for, but it is unclear whether such a survey-based sample would be representative of the entire hedge fund industry. Also, it is not clear whether hedge fund managers have an incentive to share this information given the competitive nature of the industry.

Our *RAI* measure has several advantages. First, by examining the marginal contribution of ChatGPT information to hedge funds' portfolio change, it is likely to pick up the usage by portfolio managers for investment analysis purposes rather than other reasons. Second, our methodology can be applied to all hedge funds with holdings information, allowing us to conduct a systematic analysis of the impact of generative AI on hedge fund performance.

Nonetheless, we recognize some limitations of *RAI*. First, our estimation is based on the signals generated by ChatGPT from earnings conference call transcripts. If a hedge fund manager uses ChatGPT to analyze data other than earnings calls, and if the signals extracted from those data are not highly correlated with the signals from the earnings data, we may get a low estimate of *RAI* even though the manager uses ChatGPT to do investment analysis. However, we are not particularly concerned about this false negative case, because it implies that our estimation of the prevalence of the adoption of generative AI by hedge funds might be too conservative. Thus, our estimates provide a lower bound for the effect.

A more concerning case is the false positive one. For instance, if a fund happens to obtain information that correlates with ChatGPT signals, but does not actually use ChatGPT, then their *RAI* may be overestimated.[25] We tackle this issue in several ways:

[25]In the hedge fund industry, many funds have in-house IT teams that build and use AI tools, including machine learning systems. However, the likelihood that these in-house tools outperform ChatGPT is low, given that ChatGPT-3.5 represents a significant technological breakthrough. Still, the in-house tools may generate signals that are correlated with ChatGPT-generated signals. Therefore, there is a concern that our measure may be confounded by the presence of in-house AI tools.

First, we utilize the introduction of ChatGPT as an exogenous shock to the availability of generative AI technology. Although such technology might be proprietarily developed by a small subset of hedge funds ahead of ChatGPT, it is not likely available to most hedge funds before the release of ChatGPT. Therefore, we can use the introduction of the ChatGPT as a shock to the availability of the technology to hedge funds. If *RAI* merely captures the cases where hedge funds use information correlated with ChatGPT signals without using ChatGPT itself, then we do not expect a significant jump in *RAI* around the introduction of the ChatGPT. In Section 3.3, we look at the time trends of *RAI* as a first validation test.

Second, in Section 3.4, we formally estimate the false positive rate using a partial F-test. The idea is that if a fund does not use generative AI, then F-test should be insignificant. If we assume no funds have access to generative AI technology prior to 2021, we can calculate the percentage of funds with a significant F-test during the earlier sample of 2016-2021 to estimate the false positive rate of *RAI*.

Third, in Section 4.1, where we examine the relation between *RAI* on hedge fund performance, we utilize a difference-in-differences (DID) setting. We examine whether *RAI*'s predictive power for fund performance significantly increases following the introduction of new GPT models based on which ChatGPT was trained. In Section 6.1, We also use the availability of ChatGPT to the general public as another DiD setting to confirm our findings.

Lastly, in Section 4.2, we leverage ChatGPT outages as exogenous shocks to the availability of ChatGPT to its users. We examine whether the effect of *RAI* on performance becomes more muted when there are major ChatGPT outages.

3.3. Time Trends in *RAI*

We explore the time trends in our *RAI* measure. This exercise serves two purposes. First, we want to detect the time-series pattern in addition to cross-sectional variation in *RAI*. Figure 2 shows the time series pattern of *RAI* of hedge fund companies from 2016 to

2023. We observe a surge in the *RAI* in 2022 and it is more pronounced in 2023.

Second, the significant increase in *RAI* in 2022 corresponds with the release of the first model in GPT 3.5 series along with API tools in March 2022. This is also consistent with the fact that ChatGPT 3.5 was later introduced in November 2022 to the public.[26] These findings serve as a validation test for our reliance measure, indicating its ability to capture the rising usage trend of ChatGPT.

[Insert Figure 2 Here]

3.4. A Formal Test of Hedge Funds' Generative AI Adoption

One important question is how widely ChatGPT is used by hedge funds. A fund is more likely to use ChatGPT if its *RAI* is higher, but how high does the *RAI* have to be in order for us to say a fund is using ChatGPT? One way to objectively determine that is a partial F-test, which formally tests whether the model's explanatory power is significantly improved by adding an additional variable. Specifically, it is calculated as

$$F_{i,t} = \frac{(RSS_{fundamental,i,t} - RSS_{AI,i,t})/p}{RSS_{AI,i,t}/(n-k)} \tag{4}$$

where $RSS_{fundamental,i,t}$ is the residual sum of squares of the model with firm fundamentals only, i.e., equation (1), while $RSS_{AI,i,t}$ is the residual sum of squares of the full model after adding the fundamental information generated by ChatGPT, i.e., equation (2). p is the number of predictors added to the full model and equals 14 in our case since we have 14 ChatGPT scores. n is the number of observations used to estimate equation (1) and (2) in a given fund quarter. k is the number of coefficients (including the intercept) in the full model and equals 20 since we have five variables about firm fundamentals, 14 ChatGPT scores, and an intercept.

[26]Sources: "New GPT-3 capabilities: Edit & insert," March 15, 2022, OpenAI; and "ChatGPT: Optimizing language models for dialogue," November 30, 2022, OpenAI.

We conduct the partial F-test for each hedge fund company-quarter. A hedge fund company is considered as a generative AI user for a quarter if its F-test is significant at the 1% level. We then calculate the percentage of funds with significant partial F-tests out of all funds for each period.

Figure 3 shows the yearly average of this percentage. The percentage was low and smooth before 2021 and increased dramatically in 2022. With a p-value of 0.01, we expect a false positive rate of 1%, so even if funds do not use generative AI at all, we will still find 1% of the funds having a positive and significant F-test. According to the figure, the percentage of funds with significant F-tests is around 3%. Thus, the false-positive rate contributed by our measure is estimated to be around 2% (ie. 3% total false-positive rate minus 1% false-positive rate contributed by the F-test itself). This estimate suggests that our measure misattributes about 2% of funds that do not use generative AI but happen to have trading strategies that correlate with ChatGPT signals as ChatGPT users.

In 2022, the percentage increases to about 21.0%, subtracting the average positive rate of 2%, we can infer that 19% of funds adopted ChatGPT in 2022. In 2023, this number remains at 18%. This is a notable adoption rate and is consistent with the speed at which the general public subscribes to ChatGPT.

[Insert Figure 3 Here]

3.5. Who are the Early Adopters?

After detecting the early movers in hedge funds, a natural question is whether larger or smaller funds are more likely to adopt generative AI first. Ex ante, it is unclear which direction the prediction takes. On the one hand, large funds are more likely to adopt AI technology quickly because they have more resources to adopt new technology. For instance, large financial companies such as BlackRock and JP Morgan have their in-house

research teams to utilize AI in investment.[27] There is anecdotal evidence that large funds already use AI for their investment decisions.[28] On the other hand, small funds have more incentives to use new tools to develop their edges and are more nimble to switch to a new technology.

To empirically examine this question, we define a hedge fund company as an *Early Adopter* if it has an insignificant F-statistic from equation (4) at the end of 2021 but a significant F-statistic at 1% level during any quarter in 2022. We link this indicator variable with hedge fund company characteristics observed at the end of 2021 and run a cross-sectional regression. Besides fund size, we also include several characteristics that are known to be important for fund performance: *Age*, *Risk*, *Turnover*, and *Past Return*.

Unlike *Size*, *Age*, or *Risk* that are either persistent over quarters or measured over the long term, *Turnover* and *Past Return* are measured within a quarter and could vary between quarters. We therefore calculate the *Average Turnover* and *Average Past Return* in 2021. Since the independent variable is a dummy variable, we use Logit, Probit, and Linear models to run the following regression:

$$Early\ Adopter_{i,2022} = \gamma \cdot FundCharacter_{i,2021} + \varepsilon_i \tag{5}$$

where $FundCharacter_{i,2021}$ include *Size*, *Age*, *Risk*, *Average Turnover*, and *Average Past Return*, and all are measured at the end of 2021.

[Insert Table 2 Here]

Table 2 shows that the coefficient on the *Size* is positive and significant, suggesting larger hedge fund companies are more likely to adopt generative AI early. Our results also show that AI adoption is not related to fund age. In addition, our findings suggest that

[27] JP Morgan's AI research program can be found here: https://www.jpmorgan.com/technology/artificial-intelligence.

[28] Sources: "JPMorgan uses quantum computing to summarize documents," December 5, 2022, Berenice Baker, *IoT World Today*; "How AI is transforming investing," Jun 15, 2023, BlackRock.

hedge funds with high turnover and low-risk portfolios are also more likely to be early adopters. Moreover, those with good past performance tend to adopt generative AI first.

4. Generative AI and Fund Performance

In this section, we test whether generative AI is associated with performance in the asset management industry with a focus on hedge fund companies first and then including other asset management firms. In addition, we use ChatGPT outages as exogenous shocks to further establish the link between *RAI* and fund performance.

4.1. *RAI* and Hedge Fund Performance

Our novel *RAI* measure captures the responsiveness of a fund manager's portfolio allocations to changes in AI-generated information. Since the prior studies show that AI-generated information is useful in predicting future corporate policies and returns (e.g.,Jha, Qian, Weber, and Yang, 2023), we hypothesize that funds with high *RAI* tend to outperform funds with low *RAI*.

To test this hypothesis, our empirical analysis starts with linking future performance and *RAI*. Since the base model for ChatGPT was released in March 2022, hedge funds likely started to use new versions of GPT models in the second quarter of 2022. Therefore, we examine the performance of hedge funds after that. We test with both raw returns and abnormal returns (i.e., alphas).

We first test the relationship between *RAI* and raw returns of hedge funds for the sample period after the introduction of GPT services, i.e., from the third quarter of 2022 to the second quarter of 2023. Specifically, we consider the following regression.

$$Return_{i,t} = \beta \cdot RAI_{i,t-1} + \gamma \cdot Control_{i,t-1} + \alpha_t + \varepsilon_{i,t} \tag{6}$$

where i and t index hedge fund investment company and quarter. *Control* includes *Size*,

Age, *Turnover*, *Risk*, and *Past Return*. α represents time (i.e., year-quarter) fixed effects.

Table 3 shows that the coefficient on *RAI* is positive and statistically significant across various specifications. This provides supportive evidence that hedge funds with higher reliance on generative AI can produce better performance in the future. The economic magnitude is also substantial. Taking the specification with time-fixed effects and control variables (Column (4)) as an example, a one-standard-deviation change in *RAI* is associated with a 0.55% increase in the quarterly portfolio return, or 2.2% annually, which is 21% of the average value of fund returns.

[Insert Table 3 Here]

Moreover, we repeat this analysis for each year between 2016 and 2023. Figure 4 suggests that the relation between *RAI* and future performance emerges in 2022 and is more pronounced in 2023, which is consistent with our prior economic intuition. The sensitivity of performance on *RAI* is not significantly different from zero prior to 2022 (except in 2019, when it is marginally significant), but is dramatically different from zero from 2022 onwards. This also helps to rule out the possibility that there are pre-existing trends prior to the ChatGPT introduction.

[Insert Figure 4 Here]

To sharpen our analysis, we consider the new development in generative AI as an exogenous shock to hedge fund investment companies and conduct a difference-in-differences (DiD) test as follows:

$$Return_{i,t} = \beta_1 \cdot RAI_{i,t-1} \times Post\ GPT_t + \beta_2 \cdot RAI_{i,t-1} + \gamma \cdot Control_{i,t-1} + \alpha_t(\beta_3 \cdot Post\ GPT_t) + \varepsilon_{i,t} \tag{7}$$

where *Post GPT* is an indicator variable equal to one if the fund performance is measured in and after the third quarter of 2022 and zero otherwise. Note that when adding time

fixed effects, α_t subsumes *Post GPT*. *Control* includes *Size*, *Age*, *Turnover*, *Risk*, and *Past Return*. The sample period for this test is from the beginning of 2016 to the second quarter of 2023. We expect to find a positive coefficient on the interaction term if generative AI has a positive effect on hedge fund performance.

[Insert Table 4 Here]

Table 4 confirms our hypothesis. The coefficient on $RAI \times Post\ GPT$ is positive and statistically significant at the 1% level across specifications with and without time-fixed effects and control variables. Again, the economic magnitude is large. Column (4) suggests that a one-standard-deviation change in *RAI* is associated with a 0.44% increase in the quarterly portfolio return, which is 17% of the average value of fund returns. Focusing on the last two columns with time fixed effects, the coefficient on *RAI* is indifferent from zero, suggesting that there is no pre-trend.

So far, we use raw portfolio return as our measure of performance. We also consider other measures for performance, including *CAPM Alpha*, *FF3 Alpha* and *FF4 Alpha*. We replace the dependent variable in equations (6) and (7) by these risk-adjusted returns and repeat these tests.

$$Alpha_{i,t} = \beta \cdot RAI_{i,t-1} + \gamma \cdot Control_{i,t-1} + \alpha_t + \varepsilon_{i,t} \tag{8}$$

$$Alpha_{i,t} = \beta_1 \cdot RAI_{i,t-1} \times Post\ GPT_t + \beta_2 \cdot RAI_{i,t-1} + \gamma \cdot Control_{i,t-1} + \alpha_t + \varepsilon_{i,t} \tag{9}$$

where i and t index hedge fund investment company and quarter. *Alpha* is *CAPM Alpha*, *FF3 Alpha* or *FF4 Alpha*. *Control* includes *Size*, *Age*, *Turnover*, *Risk*, and *Past Return*. α represents time (i.e., year-quarter) fixed effects.

[Insert Table 5 Here]

We report the results in Table 5. Our findings still hold when using risk-adjusted returns as dependent variables. The economic magnitude is also non-trivial, an interquartile increase in *RAI* is associated with a 70-bps to 115-bps increase in quarterly risk-adjusted returns, depending on factor models used, and equivalent to 2.8%-4.6% annual alphas.

4.2. ChatGPT Outages

To provide further support for the effect of AI on the fund industry, we use ChatGPT outages as exogenous shocks. We hypothesize that if the effect of *RAI* on fund performance is indeed from ChatGPT, this effect will be smaller when there are major ChatGPT outages because fund managers cannot use ChatGPT to aid their decisions when the tool is down.

To test this idea, we collect outage occurrences from the OpenAI website. From 2023 to the first quarter of 2024, there were a total of 42 outages, 13 of which lasted more than one hour. To exploit outages as exogenous shocks on the usage of ChatGPT, we estimate the following DiD regression with a focus on risk-adjusted returns,

$$Alpha_{i,t} = \beta_1 \cdot RAI_{i,t-1} \times Post\ GPT_t \times\ Outage_t + \beta_2 \cdot RAI_{i,t-1} \times Post\ GPT_t + \beta_3 \cdot RAI_{i,t-1} + \gamma \cdot Control_{i,t-1} + \alpha_t + \varepsilon_{i,t}, \quad (10)$$

where *Outage* is an indicator variable equal to one during the quarters with higher than median outage occurrences. 2023Q1 and 2023Q3 had a larger number of outages than the median number (9 outages) of outages during the sample period. i and t index hedge fund investment company and quarter. *Alpha* is *CAPM Alpha*, *FF3 Alpha* or *FF4 Alpha*. *Control* includes *Size*, *Age*, *Turnover*, *Risk*, and *Past Return*. α represents time (i.e., year-quarter) fixed effects.

[Insert Table 6 Here]

We report the results in Table 6. The coefficient on the triple interaction term is negative

and significant, suggesting that the effect of ChatGPT on fund performance is much weaker when ChatGPT experiences a large number of outages. The economic magnitude is large. Column (2) suggests that the ChatGPT effect on fund performance is 64% smaller during quarters with many outages. The reduction in the effect is even larger (78% to 80%) when we use *FF3 Alpha* or *FF4 Alpha*. While this magnitude seems to be large, it is reasonable because ChatGPT outages may negatively affect fund managers not only because the tool is unavailable during those times, but also because frequent outages may undermine managers' confidence in the tool and reduce their likelihood of using it outside of the outage windows.

Overall, the result from ChatGPT outages provides further support for the effect of ChatGPT adoption and usage on hedge fund performance because it suggests that the effect is less likely due to other driving forces. This result also indicates a causal relationship between AI adoption and fund performance.

4.3. *RAI* and Other Asset Management Firms

Hedge funds are arguably pioneers in applying AI and machine learning to their investment strategies. However, recent developments in intuitive AI applications such as ChatGPT make it more accessible to broader user groups such as mutual funds and other money managers. Therefore, we examine whether these asset managers also use generative AI and, more importantly, boost their portfolio performance.

We reproduce the analyses from equations (6) and (7) in Section 4.1 for asset management companies that do not operate hedge funds (which we label as *Non-Hedge Funds*) and report the findings in Table 7.

[Insert Table 7 Here]

In contrast to their hedge fund peers, these asset managers cannot generate superior performance, albeit having an increasing usage of generative AI after it becomes available.

This finding holds for *Post GPT* period as well as using a DiD setting. In addition, Table 7 Panel B presents a direct comparison between these non-hedge funds and hedge funds. We observe a clear difference between these two groups of investment companies. Columns (1) and (2) show that, during *Post GPT* period, only hedge funds can boost their performance by using generative AI. Columns (3) and (4) further confirm that the advent of generative AI allows hedge funds to transform AI applications into better performance.

These findings could be due to various advantages that hedge funds enjoy relative to other asset managers. For example, hedge funds have better access to data, and can process data and execute trades more quickly. They may also combine other information or trading skills with AI-generated information to further improve performance. In sum, the results indicate that generative AI may need to be combined with other resources or expertise in order to produce superior investment returns.

4.4. AI Disparity Among Funds

One important question is which firms benefit more from the use of generative AI. One hypothesis is that this convenient and powerful new tool would help to level the playing field for hedge fund companies with different resources and capabilities. On the other hand, large hedge funds may be able to combine their resources with generative AI to further increase their competitive advantage. Therefore, it is an empirical question which types of hedge fund companies are most effective in utilizing generative AI in their investment.

In Figure 5, we compare the sensitivity of fund company performance to *RAI*, i.e., the coefficient of the interaction term in equation (7), for top and bottom quintile hedge fund companies defined by firm characteristics, including *Size*, *Age*, *Risk*, *Turnover*, and *Past Return*. The results show that larger, older, and more active fund companies are able to leverage generative AI and generate superior performance, while the usage by small, younger, and less active firms does not yield significant returns. Sorting on other

characteristics does not generate significant differences. Overall, the results suggest that fund companies with more resources benefit the most from the use of generative AI, suggesting a potential synergy between the novel AI tool with other resources such as data and expertise, consistent with what we find in Section 3.5 and 4.3.

[Insert Figure 5 Here]

5. How does AI Help Hedge Fund Performance?

So far, we show that generative AI helps hedge funds obtain better performance. In this section, we explore potential economic channels. In particular, we test two channels. First, we examine whether hedge funds invest more in human capital in AI so that they can use the tools better. Second, we examine whether generative AI helps funds to analyze certain data better.

5.1. Combination with AI Talent

To understand the *RAI* effect, we explore one potential channel of AI investment by hedge funds. Anecdotal evidence shows that hedge funds heavily invest in human capital in the area of AI so that they can have the talent to use the tools better. To test this idea, we focus on a subset of hedge funds with greater capacity in applying AI tools and expect our findings to be more pronounced. Following Cao, Jiang, Yang, and Zhang (2023), we classify hedge funds that have employed AI-skilled workers as *AI Hedge Fund* and hypothesize that these funds have a greater likelihood of using generative AI to produce a better performance.

[Insert Table 8 Here]

Table 8 shows that our findings hold for all hedge funds and, more importantly, are much stronger within AI hedge funds. With respect to economic magnitude, among all

hedge funds, a one-standard-deviation increase in *RAI* leads to an increase of 0.39% in quarterly portfolio return. On top of that, it results in a significant increase of 1.65% – adding up to a total increase of 2.04% in quarterly return – for AI hedge funds. These results suggest that the combination of AI talent with the tools is likely to be a driving force for *RAI* effect on fund performance. This is consistent with the complementarity between humans and machines documented in the existing literature (e.g., Cao, Jiang, Wang, and Yang, 2024).

5.2. Strength of Analyzing Certain Data

Another potential channel is that ChatGPT is good at analyzing certain data and providing predictions. To test this idea, we further explore the granular components of AI-generated information. The 14 GPT scores generated by ChatGPT with earnings conference calls can be naturally separated into three groups: 1) Macro, 2) Firm Policy, and 3) Firm Performance.[29] We repeat our methodology for defining *RAI*, and every time only add information generated by ChatGPT for each respective group in equation (2). We then create the decomposed *RAI* measures for each group: *RAI Macro*, *RAI Firm Policy*, and *RAI Firm Performance*, helping us pinpoint what kind of information hedge funds use to provide superior performance.

[Insert Table 9 Here]

We repeat regression analyses in equations (7) and (9) by replacing *RAI* with one of the three decomposed measures and report the results in Table 9. We observe that the interaction between *RAI Macro* and *Post GPT* is indistinguishable from zero, suggesting that reliance on AI information about macroeconomics does not help with fund performance.

[29]The Macro group contains information regarding the global economy, the US economy and a firm's industry; the Firm Policy group pertains to a firm's wages, employment, capital expenditure, and cost of capital; the Firm Performance group is about a firm's earnings, revenue, financial prospects, and product market.

In contrast, both *RAI Firm Policy* and *RAI Firm Performance* have a significant and positive relation with hedge fund performance during *Post GPT* period. AI-generated information about firm policy is particularly useful as the magnitude of the coefficient is more than twice as much as that on AI information about firm performance.

These findings suggest possible channels that generative AI tools enhance performance in asset management. First, generative AI is more useful for hedge funds to select individual stocks rather than conduct sector or market timing conditional on the macroeconomy. One notable advantage of generative AI is that it can process a tremendous amount of textual data and is especially efficient when hedge fund companies face thousands of stocks to make informed investment decisions. On the other hand, generative AI is less important when hedge funds need information about the industry, U.S. market, or global market since they are unlikely to look into portfolio firms' filings or conference calls to collect such information. Moreover, our findings also indicate that firm policy is informative about stock return, consistent with Jha, Qian, Weber, and Yang (2023). Therefore, generative AI helps hedge funds extract valuable information from voluminous public data and reap benefits from the stock market.

6. Robustness

In this section, we provide two robustness tests. First, we consider the alternative date for the DiD test. Second, we construct an alternative measure of *RAI*.

6.1. An Alternative DiD Test

In the main specification, we use the release of the base GPT 3.5 model in March 2022 as the cutoff for the DiD analysis. The refined model was released as ChatGPT 3.5 through the chat-based interface in November 2022 to the public. Therefore, an alternative way to conduct the DiD analysis is to use the formal release date of ChatGPT 3.5 to define the

post-period.

We construct a dummy variable *Post ChatGPT*, which equals one for performance in the first quarter of 2023 and onwards, and zero otherwise. We then re-run the performance regression and report the results in Table 10. The coefficient on the interaction term is positive and significant, suggesting that funds that have a higher *RAI* tend to outperform after the release of ChatGPT. This is consistent with our main specification.

[Insert Table 10 Here]

6.2. An Alternative Measure of *RAI*

In addition, we consider an alternative measure for our key variable of interest. As a robustness check, we create *Alt. RAI*, defined as the percentage increase in R^2, i.e.,

$$Alt.\ RAI_{i,t} = \frac{R^2_{AI,i,t} - R^2_{fundamental,i,t}}{R^2_{fundamental,i,t}}. \tag{11}$$

The rationale for this alternative measure is to benchmark against the explanatory power of fundamental information. We then redo our analyses in Tables 3, 4, and 5. Our results, reported in Table 11 are qualitatively similar when we use this alternative measure and show again that the adoption of generative AI is associated with significant increases in hedge fund performance, both in terms of raw and risk-adjusted returns.

[Insert Table 11 Here]

7. Conclusion

In this paper, we develop a novel measure of the usage or reliance on generative AI (*RAI*) of investment companies based on their portfolio holdings and AI-predicted information. We study the adoption and implications of generative AI in hedge funds and

other asset management companies. Utilizing *RAI*, we find a dramatic increase in the use of generative AI by hedge fund companies after the introduction of ChatGPT.

Hedge fund companies with higher *RAI* produce superior returns, both unadjusted and risk-adjusted. For example, an interquartile change in the *RAI* is associated with an increase of 3 to 5% in annualized hedge fund returns. In a triple difference-in-differences test, we exploit ChatGPT outages as exogenous shocks and show that the effect of generative AI adoption on performance is substantially reduced with major outages.

In investigating the source of the superior performance, we find hedge fund companies generate more returns from using AI-predicted firm-specific information related to firm policies and performance than from macroeconomic and sectorwise information. Not all investment companies benefit equally from the invention of generative AI: Non-hedge fund companies do not produce significant returns. Furthermore, large and more active hedge fund companies adopt the technology early and perform better than others.

Overall, our findings shed light on the use and implications of generative AI technology and suggest that despite being intuitive to use, generative AI may need to be combined with other resources, such as data and expertise, to be productive for the adopting companies. Importantly, the benefits of generative AI predominantly accrue to larger players who possess the resources to effectively implement and leverage such technologies, potentially widening disparities within the industry. Our findings also carry implications in broader societal contexts, as the increasingly wide adoption of AI[30] has the potential to not only increase productivity but also exacerbate inequality.[31]

[30]See, for example, "JPMorgan pitches in-house chatbot as AI-based research analyst," July 26, 2024, Stephen Morris and Joshua Franklin, *Financial Times*.

[31]This echoes recent debates about the effects of AI, e.g., Acemoglu and Johnson (2024), "Unregulated AI Will Worsen Inequality, Warns Nobel-Winning Economist Joseph Stiglitz," August 1, 2023, Sophie Bushwick, *Scientific American*. "AI's economic peril to democracy," March 14, 2024, Stephanie A. Bell and Anton Korinek, *Brookings*.

References

Abis, Simona, and Laura Veldkamp, 2024, The changing economics of knowledge production, *Review of Financial Studies* 37, 89–118.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial intelligence and jobs: Evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.

Acemoglu, Daron, and Simon Johnson, 2024, Learning from Ricardo and Thompson: Machinery and labor in the early industrial revolution and in the age of artificial intelligence, *Annual Review of Economics* 16, 597–621.

Acikalin, Utku, Tolga Caskurlu, Gerard Hoberg, and Gordon M. Phillips, 2022, Intellectual Property Protection Lost and Competition: An Examination Using Large Language Models, *Working paper, Available on SSRN 4023622* .

Agarwal, Vikas, Wei Jiang, Yuehua Tang, and Baozhong Yang, 2013, Uncovering hedge fund skill from the portfolio holdings they hide, *Journal of Finance* 68, 739–783.

Agarwal, Vikas, Stefan Ruenzi, and Florian Weigert, 2023, Unobserved performance of hedge funds, *Journal of Finance, forthcoming* .

Aragon, George O, Michael Hertzel, and Zhen Shi, 2013, Why do hedge funds avoid disclosure? evidence from confidential 13f filings, *Journal of Financial and Quantitative Analysis* 48, 1499–1518.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, 2023, Firm investments in artificial intelligence technologies and changes in workforce composition, *NBER Volume on Technology, Productivity, and Economic Growth, forthcoming* .

Bai, John Jianqiu, Nicole M Boyson, Yi Cao, Miao Liu, and Chi Wan, 2023, Executives vs. chatbots: Unmasking insights through human-AI differences in earnings conference Q&A, *Working paper, Available at SSRN 4480056* .

Bonelli, Maxime, 2023, Data-driven investors, *Working paper, HEC Paris* .

Bonelli, Maxime, and Thierry Foucault, 2023, Displaced by big data: Evidence from active fund managers, *Working paper, Available at SSRN 4527672* .

Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond, 2023, Generative AI at work, *Working paper, National Bureau of Economic Research* .

Cao, Charles, Yong Chen, Bing Liang, and Andrew Lo, 2013, Can hedge funds time market liquidity?, *Journal of Financial Economics* 109, 493–516.

Cao, Sean, Wei Jiang, Junbo Wang, and Baozhong Yang, 2024, From man vs. machine to man+ machine: The art and AI of stock analyses, *Journal of Financial Economics* 160, 103910.

Cao, Sean, Wei Jiang, Baozhong Yang, and Alan L Zhang, 2023, How to talk when a machine is listening: Corporate disclosure in the age of AI, *Review of Financial Studies* 36, 3603–3642.

Carhart, Mark M, 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.

Chen, Huaizhi, Lauren Cohen, Umit Gurun, Dong Lou, and Christopher Malloy, 2020, IQ from IP: Simplifying search in portfolio choice, *Journal of Financial Economics* 138, 118–137.

Chen, Yong, Bing Han, and Jing Pan, 2021, Sentiment trading and hedge fund returns, *Journal of Finance* 76, 2001–2033.

Chen, Yong, and Bing Liang, 2007, Do market timing hedge funds time the market?, *Journal of Financial and Quantitative Analysis* 42, 827–856.

Colliard, Jean-Edouard, Thierry Foucault, and Stefano Lovo, 2022, Algorithmic pricing and liquidity in securities markets, *HEC Paris Research Paper No. FIN-2022-1459* .

Cong, Lin William, Ke Tang, Jingyuan Wang, and Yang Zhang, 2021, Alphaportfolio: Direct construction through deep reinforcement learning and interpretable AI, *Working paper, Available at SSRN 3554486* .

Dou, Winston Wei, Itay Goldstein, and Yan Ji, 2023, AI-powered trading, algorithmic collusion, and price efficiency, *Working paper, Available at SSRN 4452704* .

Dugast, Jérôme, and Thierry Foucault, 2023, Equilibrium data mining and data abundance, *Journal of Finance, forthcoming* .

D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G Rossi, 2019, The promises and pitfalls of robo-advising, *Review of Financial Studies* 32, 1983–2020.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, 2023, Generative AI and firm values, *Working paper, National Bureau of Economic Research* .

Fama, Eugene F, and Kenneth R French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fisher, Adlai, Charles Martineau, and Jinfei Sheng, 2022, Macroeconomic attention and announcement risk premia, *Review of Financial Studies* 35, 5057–5093.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, 2022, Predictably unequal? the effects of machine learning on credit markets, *Journal of Finance* 77, 5–47.

Garcia, Diego, Xiaowen Hu, and Maximilian Rohrer, 2023, The colour of finance words, *Journal of Financial Economics* 147, 525–549.

Greene, William H., 2002, *Econometric Analysis* (Prentice Hall).

Greig, Fiona, Tarun Ramadorai, Alberto G Rossi, Stephen P Utkus, and Ansgar Walther, 2022, Algorithm aversion: Theory and evidence from robo-advice, *Working paper, Available at SSRN 4301514* .

Griffin, John M, and Jin Xu, 2009, How smart are the smart guys? a unique view from hedge fund stock holdings, *Review of Financial Studies* 22, 2531–2570.

Grossman, Sanford, and Joseph Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.

Hansen, Anne Lundgaard, and Sophia Kazinnik, 2023, Can ChatGPT decipher Fedspeak?, *Working paper, Available at SSRN 4399406* .

Hoberg, Gerard, and Gordon Phillips, 2010, Product market synergies and competition in mergers and acquisitions: A text-based analysis, *Review of Financial Studies* 23, 3773–3811.

Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of Political Economy* 124, 1423–1465.

Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang, 2023, ChatGPT and corporate policies, *Working paper, National Bureau of Economic Research* .

Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang, 2024, Harnessing generative AI for economic insights, *Working paper* .

Jiang, Jingwen, Bryan T Kelly, and Dacheng Xiu, 2022, Expected returns and large language models, *Working paper, Available at SSRN 4416687* .

Kacperczyk, Marcin, and Amit Seru, 2007, Fund manager use of public information: New evidence on managerial skills, *Journal of Finance* 62, 485–528.

Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng, 2005, On the industry concentration of actively managed equity mutual funds, *Journal of Finance* 60, 1983–2011.

Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng, 2008, Unobserved actions of mutual funds, *Review of Financial Studies* 21, 2379–2416.

Kim, Alex G, Maximilian Muhn, and Valeri V Nikolaev, 2023, Bloated disclosures: Can ChatGPT help investors process information?, *Chicago Booth Research Paper* .

Korinek, Anton, 2023, Generative AI for economic research: Use cases and implications for economists, *Journal of Economic Literature* 61, 1281–1317.

Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan, 2021, Measuring corporate culture using machine learning, *Review of Financial Studies* 34, 3265–3315.

Li, Kai, Feng Mai, Rui Shen, Chelsea Yang, and Tengfei Zhang, 2023, Dissecting corporate culture using generative AI – Insights from analyst reports, *Working paper, Available at SSRN 4558295* .

Lopez-Lira, Alejandro, and Yuehua Tang, 2023, Can ChatGPT forecast stock price movements? return predictability and large language models, *Working paper, Available at SSRN 4412788* .

Loughran, Tim, and Bill McDonald, 2011, When is a liability not a liability? textual analysis, dictionaries, and 10-ks, *Journal of Finance* 66, 35–65.

Lyonnet, Victor, and Léa H Stern, 2022, Venture capital (mis) allocation in the age of AI, *Working paper, Ohio State University* .

Martin, Ian, and Stefan Nagel, 2022, Market efficiency in the age of big data, *Journal of Financial Economics* 145, 154–177.

Rossi, Alberto G, and Stephen P Utkus, 2024, Who benefits from robo-advising? evidence from machine learning, *Journal of Financial Economics, forthcoming* .

Sialm, Clemens, Zheng Sun, and Lu Zheng, 2020, Home bias and local contagion: Evidence from funds of hedge funds, *Review of Financial Studies* 33, 4771–4810.

Sun, Zheng, Ashley Wang, and Lu Zheng, 2012, The road less traveled: Strategy distinctiveness and hedge fund performance, *Review of Financial Studies* 25, 96–143.

Tetlock, Paul C, 2007, Giving content to investor sentiment: The role of media in the stock market, *Journal of Finance* 62, 1139–1168.

Titman, Sheridan, and Cristian Tiu, 2011, Do the best hedge funds hedge?, *Review of Financial Studies* 24, 123–168.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, 2017, Attention is all you need, *Advances in neural information processing systems* 30.

Webb, Michael, 2019, The impact of artificial intelligence on the labor market, *Working Paper, Available at SSRN 3482150* .

Wermers, Russ, 2000, Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses, *Journal of Finance* 55, 1655–1695.

Figure 1. Timeline of ChatGPT

This figure presents the timeline of the milestones in the development of ChatGPT.

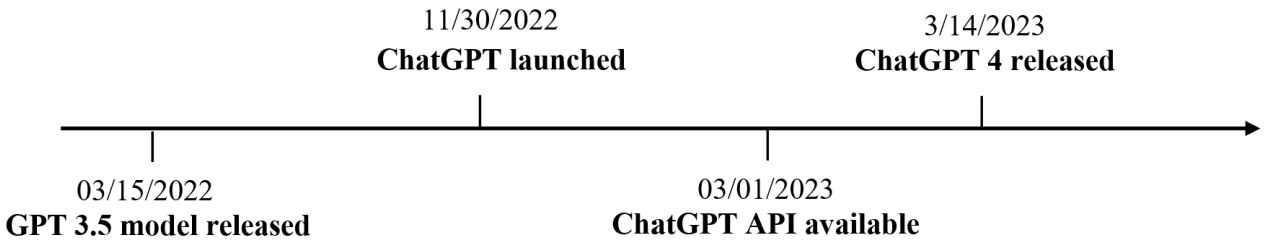


Figure 2. Trend of RAI

This figure plots the average *RAI* from 2016 to 2023. *RAI* measures the extent to which AI-generated information influences trades of hedge fund companies, defined in Table A1 of the Appendix.

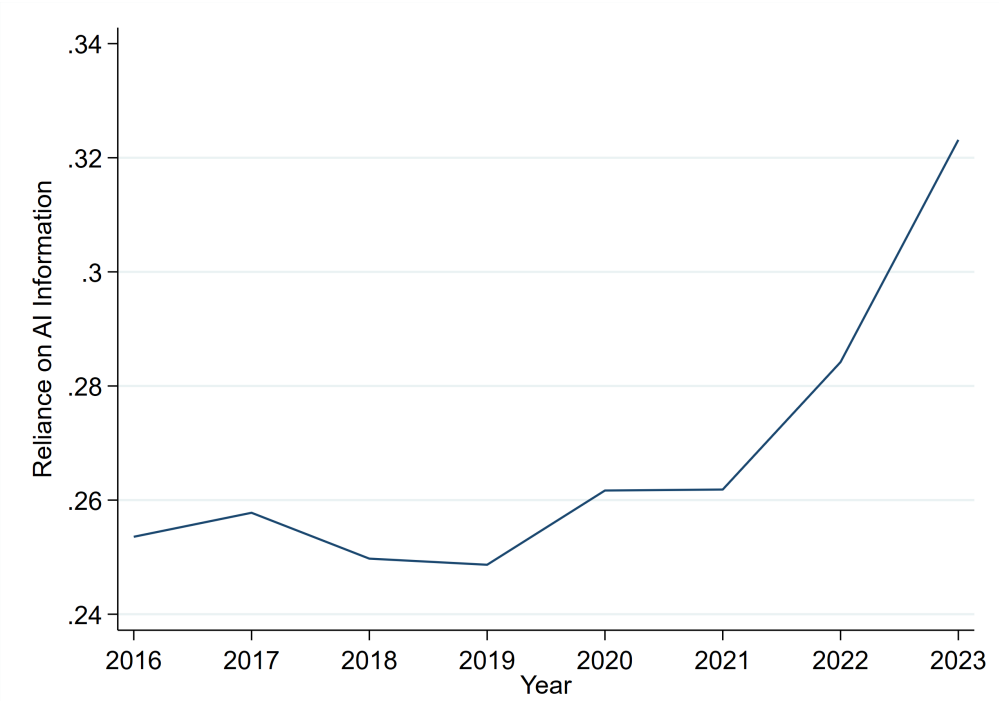


Figure 3. Trend of Generative AI Adoptiom

This figure plots the generative AI adoption from 2016 to 2023.

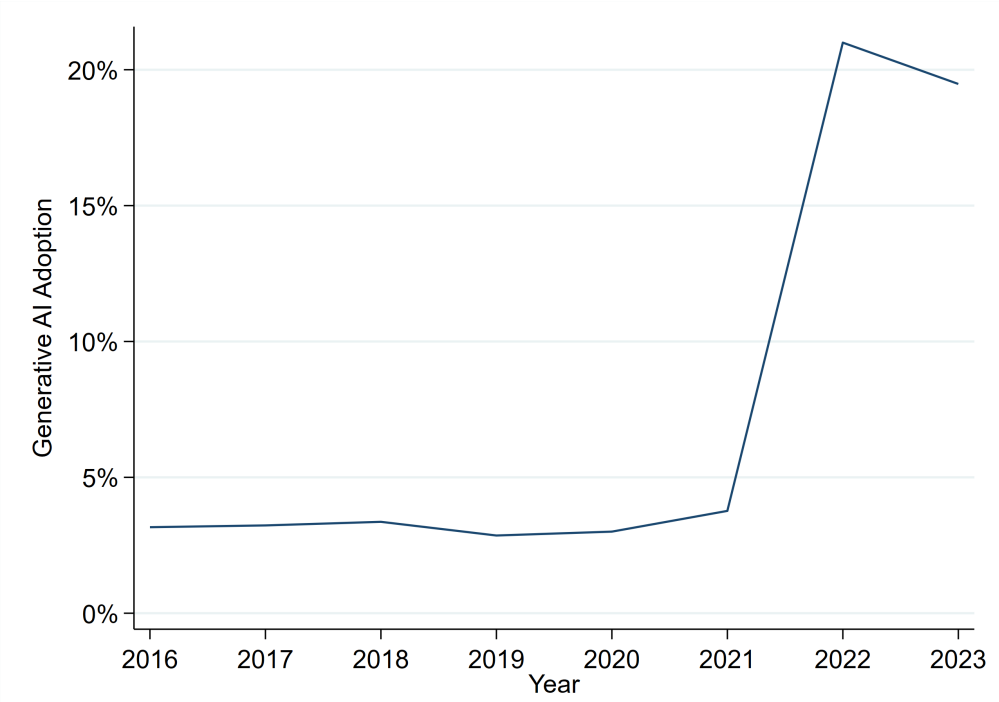


Figure 4. Trend of Performance-RAI Sensitivity

This figure plots the coefficient of *RAI* and its 95% confidence interval in the following regression for each year between 2016 and 2023:

$$Return_{i,t} = \beta \cdot RAI_{i,t-1} + \gamma \cdot Control_{i,t-1} + \varepsilon_{i,t}$$

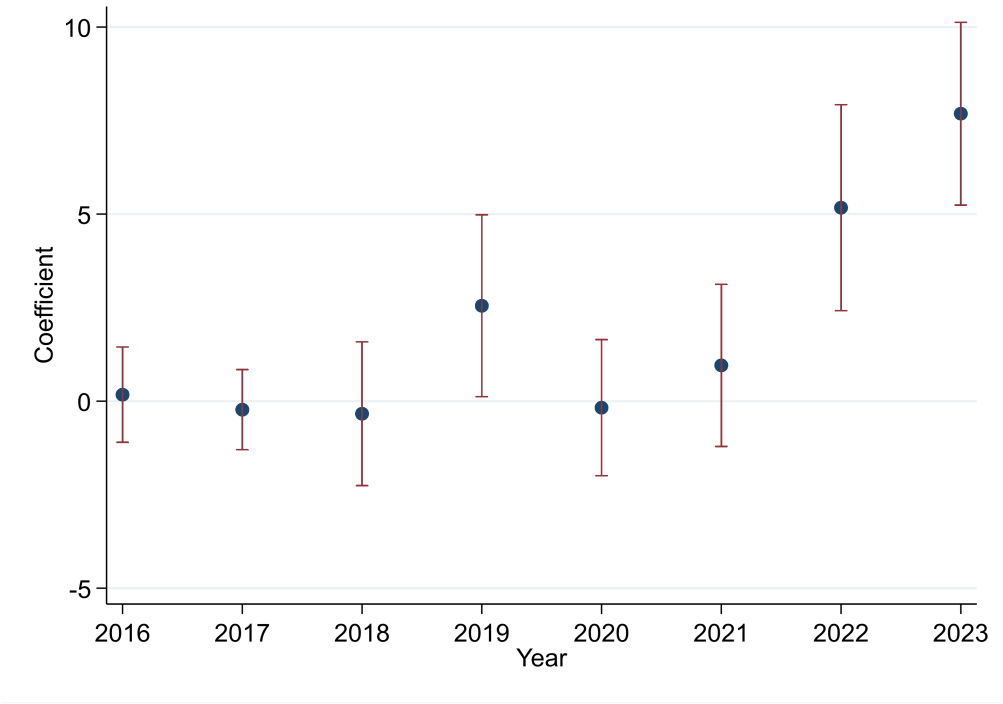
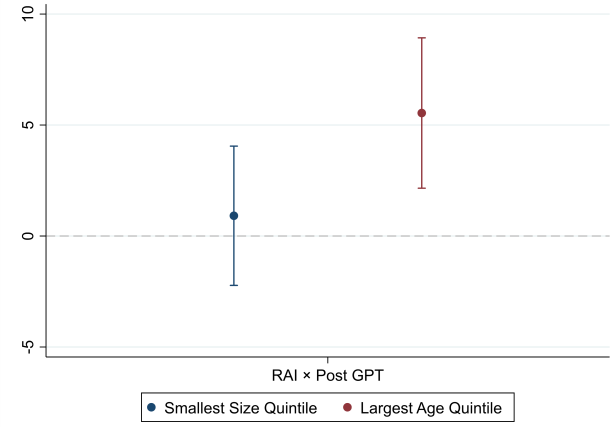
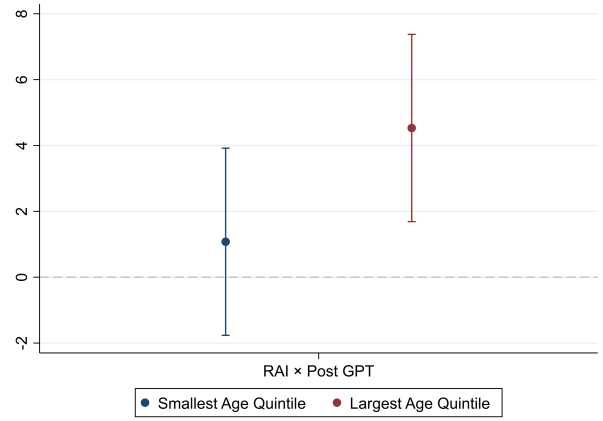


Figure 5. RAI and Hedge Fund Company Characteristics

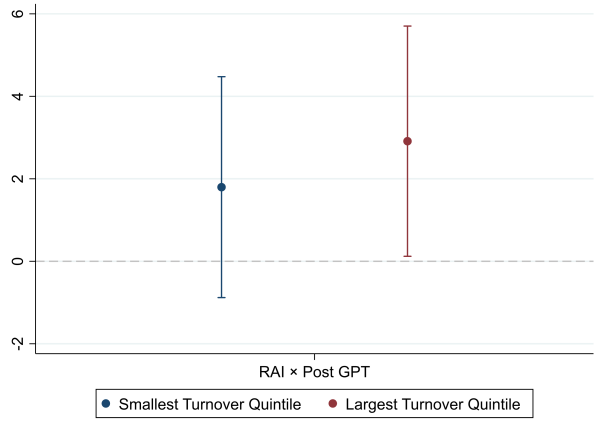
This figure plots the coefficient of interaction term for subsamples partitioned by fund company characteristics: $Return_{i,t} = \beta_1 \cdot RAI_{i,t-1} \times Post\ GPT_t + \beta_2 \cdot RAI_{i,t-1} + \gamma \cdot Control_{i,t-1} + \alpha_t + \varepsilon_{i,t}$



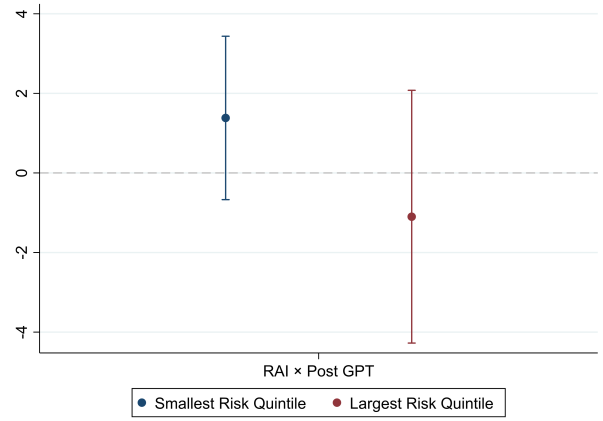
(a) Size Quintiles



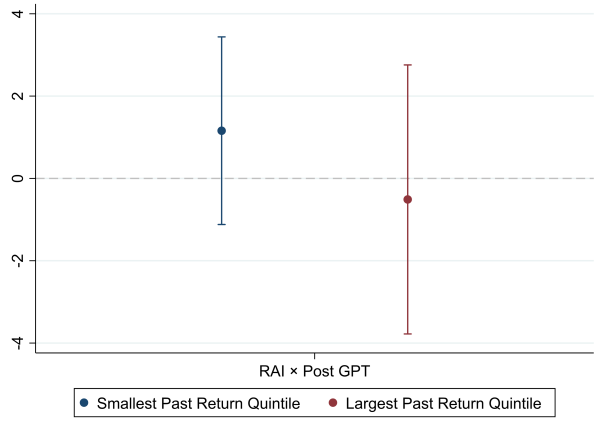
(b) Age Quintiles



(c) Turnover Quintiles



(d) Risk Quintiles



(e) Past Return Quintiles

Table 1. Summary Statistics

This table reports the summary statistics. *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Return* is the portfolio holdings return, expressed in percentage points (%). *CAPM Alpha* (*FF3 Alpha*/*FF4 Alpha*) is the portfolio holdings return, after adjusting for the market risk factor (Fama-French three factors/Fama-French-Carhart four factors), expressed in percentage points (%). *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Turnover* is the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years.

	(1)	(2)	(3)	(4)	(5)	(6)
Variables	N	Mean	St. Dev.	P25	Median	P75
RAI	10,762	0.260	0.216	0.079	0.203	0.399
Alt. RAI	10,762	3.631	4.606	1.145	2.136	4.155
Return	10,762	2.638	10.750	-1.623	3.736	7.837
CAPM Alpha	10,762	-1.516	4.563	-3.401	-1.296	0.487
FF3 Alpha	10,762	-1.640	3.793	-3.196	-1.354	0.199
FF4 Alpha	10,762	-1.676	3.889	-3.240	-1.392	0.161
Size	10,762	7.019	1.642	5.831	6.855	8.014
Age	10,762	16.260	8.739	9.250	15.000	21.500
Turnover	10,762	0.175	0.157	0.053	0.119	0.261
Risk	10,762	0.094	0.055	0.049	0.084	0.128

Table 2. Characteristics of Early Generative AI Adopters

This table reports the characteristics of early adopters of generative AI. *Early Adopter* is an indicator variable equal to one if a hedge fund company has an insignificant F-statistic from equation (4) at the end of 2021 but a significant F-statistic at 1% level in 2022. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Average Turnover* is the annual average of *Turnover*, the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Average Past Return* is the annual average of quarterly portfolio holdings return. These characteristics are calculated at the end of December 2021. The *t*-statistics, in parentheses, are based on standard errors adjusted for heteroskedasticity. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)
Model	Logit	Probit	Linear
Dep. Var.	*Early Adopter*		
Size	0.436***	0.263***	0.091***
	(5.74)	(5.83)	(6.80)
Age	-0.006	-0.004	-0.001
	(-0.44)	(-0.41)	(-0.21)
Average Turnover	1.575*	0.917*	0.335*
	(1.93)	(1.84)	(1.94)
Risk	-5.405**	-3.199**	-1.059**
	(-2.17)	(-2.10)	(-2.11)
Average Past Return	0.125**	0.069***	0.023***
	(2.49)	(2.66)	(2.79)
Observations	372	372	372
Pseudo R-squared	0.116	0.114	
R-squared			0.144

Table 3. RAI and Hedge Fund Performance

This table reports the relation between performance and reliance on AI information. *Return* is the portfolio holdings return. *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Turnover* is the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Past Return* is the lagged *Return*. The sample period is from the third quarter of 2022 to the second quarter of 2023. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)	(4)
Dep. Var.		*Return*		
RAI	5.549***	4.825***	2.333***	2.545***
	(7.90)	(5.79)	(3.85)	(3.44)
Size		0.086		0.077
		(0.89)		(0.80)
Age		-0.004		-0.010
		(-0.24)		(-0.59)
Turnover		3.796***		1.649
		(3.14)		(1.59)
Risk		10.109*		2.239
		(1.77)		(0.44)
Past Return		0.135***		-0.192***
		(7.37)		(-4.99)
Observations	1,001	1,001	1,001	1,001
R-squared	0.035	0.096	0.430	0.454
Time FE	No	No	Yes	Yes

Table 4. RAI and Hedge Fund Performance: DiD

This table reports the relation between performance and reliance on AI information. *Return* is the portfolio holdings return. *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Post GPT* is an indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Turnover* is the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Past Return* is the lagged *Return*. The sample period is from the first quarter of 2016 to the second quarter of 2023. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)	(4)
Dep. Var.		*Return*		
RAI × *Post GPT*	4.557***	6.064***	1.949***	2.035***
	(5.63)	(6.64)	(2.95)	(3.10)
RAI	0.992**	-0.518	0.384	-0.183
	(2.32)	(-1.10)	(1.40)	(-0.66)
Size		-0.193***		-0.093***
		(-3.81)		(-3.59)
Age		-0.012		0.013***
		(-1.44)		(2.80)
Turnover		2.119***		0.261
		(4.14)		(1.07)
Risk		28.275***		9.205***
		(12.87)		(4.23)
Past Return		-0.183***		0.101***
		(-16.45)		(5.92)
Post GPT	-4.101***	-5.882***		
	(-16.03)	(-18.50)		
Observations	10,762	10,762	10,762	10,762
R-squared	0.007	0.056	0.795	0.798
Time FE	No	No	Yes	Yes

Table 5. RAI and Hedge Fund Performance: Risk-Adjusted Returns

This table reports the relation between performance and reliance on AI information. *CAPM Alpha* (*FF3 Alpha*/*FF4 Alpha*) is the portfolio holdings return, after adjusting for the market risk factor (Fama-French three factors/Fama-French-Carhart four factors). *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Post GPT* is an indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Turnover* is the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Past Return* is the lagged *Return*. The sample period is from the third quarter of 2022 to the second quarter of 2023 in Panel A and from the first quarter of 2016 to the second quarter of 2023 in Panel B. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

Panel A: During Post-GPT period

	(1)	(2)	(3)
Dep. Var.	*CAPM Alpha*	*FF3 Alpha*	*FF4 Alpha*
RAI	3.587***	2.636***	2.122***
	(5.43)	(3.83)	(2.86)
Size	0.235**	0.047	0.021
	(2.40)	(0.55)	(0.24)
Age	-0.028*	-0.016	-0.014
	(-1.75)	(-1.02)	(-0.86)
Turnover	1.559	1.508	1.137
	(1.40)	(1.40)	(1.06)
Risk	0.177	-2.138	-0.707
	(0.04)	(-0.51)	(-0.16)
Past Return	-0.229***	-0.031	-0.035
	(-6.51)	(-0.84)	(-0.94)
Observations	1,001	1,001	1,001
R-squared	0.127	0.040	0.039
Time FE	Yes	Yes	Yes

Panel B: DiD

	(1)	(2)	(3)
Dep. Var.	*CAPM Alpha*	*FF3 Alpha*	*FF4 Alpha*
RAI × *Post GPT*	3.298***	2.929***	2.538***
	(4.94)	(4.78)	(3.96)
RAI	-0.334	-0.411*	-0.399
	(-1.20)	(-1.67)	(-1.58)
Size	-0.116***	-0.114***	-0.119***
	(-3.23)	(-3.69)	(-3.89)
Age	0.001	-0.001	-0.001
	(0.19)	(-0.23)	(-0.22)
Turnover	-0.477	-0.657**	-0.660**
	(-1.50)	(-2.28)	(-2.30)
Risk	-6.729**	-14.062***	-13.692***
	(-2.51)	(-6.89)	(-6.27)
Past Return	0.042***	-0.005	-0.008
	(3.25)	(-0.52)	(-0.76)
Observations	10,762	10,762	10,762
R-squared	0.108	0.094	0.090
Time FE	Yes	Yes	Yes

Table 6. RAI and Hedge Fund Performance: ChatGPT Outages

This table reports how ChatGPT outage affects the relation between performance and reliance on AI information. *Return* is the portfolio holdings return. *CAPM Alpha* (*FF3 Alpha*/*FF4 Alpha*) is the portfolio holdings return, after adjusting for the market risk factor (Fama-French three factors/Fama-French-Carhart four factors). *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Post GPT* is an indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise. *Outage* is an indicator variable equal to one during the quarters with higher than median outage occurrences, and zero otherwise. Control variables include *Size*, *Age*, *Turnover*, *Risk*, and *Past Return*, defined in Table A1 of the Appendix. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)	(4)
Dep. Var.	*Return*	*CAPM Alpha*	*FF3 Alpha*	*FF4 Alpha*
RAI × *Post GPT* × *Outage*	-0.219	-2.332**	-2.569**	-2.186**
	(-0.18)	(-2.02)	(-2.38)	(-1.98)
RAI × *Post GPT*	2.060***	3.570***	3.229***	2.792***
	(2.87)	(4.89)	(4.81)	(3.97)
Observations	10,762	10,762	10,762	10,762
R-squared	0.798	0.108	0.095	0.090
Control variables	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes

Table 7. RAI: Hedge Funds vs Other Asset Management Firms

This table reports the relation between performance and reliance on AI information. *Return* is the portfolio holdings return. *FF4 Alpha* is the portfolio holdings return after adjusting for Fama-French-Carhart four factors. *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Post GPT* is an indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise. *Hedge Fund* is an indicator variable if an investment company is a hedge fund company and zero otherwise. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Turnover* is the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Past Return* is the lagged *Return*. In both panels, the sample period is from the third quarter of 2022 to the second quarter of 2023 for columns (1) to (2) and from the first quarter of 2016 to the second quarter of 2023 for columns (3) to (4). The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

Panel A: Non-hedge fund companies

	(1)	(2)	(3)	(4)
Sampe Period	Post-GPT period		DiD	
Dep. Var.	*Return*	*FF4 Alpha*	*Return*	*FF4 Alpha*
RAI × *Post GPT*			-0.235	0.155
			(-0.64)	(0.52)
RAI	-0.127	0.447	-0.127	0.357***
	(-0.32)	(1.49)	(-1.23)	(3.27)
Size	0.021	-0.025	-0.031***	-0.055***
	(0.50)	(-0.81)	(-2.97)	(-4.11)
Age	0.011*	0.006	0.010***	-0.005**
	(1.69)	(1.17)	(6.68)	(-2.08)
Turnover	1.075	0.205	-0.840***	-2.058***
	(1.16)	(0.23)	(-3.82)	(-6.42)
Risk	18.062***	2.445	16.802***	-15.488***
	(3.83)	(0.68)	(12.01)	(-9.93)
Past Return	-0.188***	-0.076***	0.081***	-0.030***
	(-7.40)	(-3.90)	(8.46)	(-5.11)
Observations	2,685	2,685	36,112	36,112
R-squared	0.711	0.077	0.885	0.148
Time FE	Yes	Yes	Yes	Yes

Panel B: Hedge fund companies vs. non-hedge fund companies

	(1)	(2)	(3)	(4)
Sampe Period	Post-GPT period		DiD	
Dep. Var.	*Return*	*FF4 Alpha*	*Return*	*FF4 Alpha*
RAI × *Post GPT* × *Hedge Fund*			2.356***	2.136***
			(3.14)	(3.30)
RAI × *Post GPT*			-0.223	0.153
			(-0.60)	(0.51)
RAI × *Hedge Fund*	2.102***	1.349**	-0.413	-0.587**
	(2.89)	(2.28)	(-1.54)	(-2.50)
RAI	0.036	0.479	-0.088	0.329***
	(0.09)	(1.59)	(-0.85)	(3.03)
Size	0.041	-0.014	-0.042***	-0.068***
	(1.04)	(-0.47)	(-4.15)	(-5.51)
Age	0.009	0.003	0.012***	-0.004**
	(1.36)	(0.49)	(7.86)	(-2.10)
Turnover	1.383**	0.730	-0.321**	-1.389***
	(1.99)	(1.13)	(-2.00)	(-6.49)
Risk	10.975***	1.626	12.596***	-15.032***
	(3.08)	(0.61)	(10.86)	(-12.32)
Past Return	-0.195***	-0.059***	0.087***	-0.021***
	(-8.88)	(-3.42)	(10.37)	(-3.97)
Observations	3,674	3,674	46,714	46,714
R-squared	0.628	0.068	0.863	0.131
Time × Company Type FE	Yes	Yes	Yes	Yes

Table 8. RAI and Hedge Fund Performance: AI Hedge Funds

This table reports how AI investment affects the relation between performance and reliance on AI information. *Return* is the portfolio holdings return. *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Post GPT* is an indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise. *AI Hedge Fund* is an indicator variable equal to one if a hedge fund has AI-skilled workers and zero otherwise. Control variables include *Size*, *Age*, *Turnover*, *Risk*, and *Past Return*, defined in Table A1 of the Appendix. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)	(4)
Dep. Var.		*Return*		
RAI × *Post GPT* × *AI Hedge Fund*	7.630***	9.288**	7.688***	9.300**
	(3.19)	(2.40)	(2.88)	(2.25)
RAI × *Post GPT*	1.793***	1.786***	1.892***	1.889***
	(2.61)	(2.59)	(2.78)	(2.76)
Observations	10,762	10,762	10,762	10,762
R-squared	0.795	0.795	0.798	0.798
Control variables	No	No	Yes	Yes
Time FE	Yes		Yes	
Time × AI Hedge Fund FE		Yes		Yes

Table 9. RAI and Hedge Fund Performance: Decomposition

This table reports how the relation between performance and reliance on AI information depends on the types of AI-generated information. *Return* is the portfolio holdings return. *FF4 Alpha* is the portfolio holdings return after adjusting for Fama-French-Carhart four factors. *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. We separate fundamental information generated by ChatGPT into three groups: 1) Macro, 2) Firm Policy, and 3) Firm Performance and create decomposed RAI measures for each respective group: *RAI Macro*; *RAI Firm Policy*; *RAI Firm Performance*. *Post GPT* is an indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise. Control variables include *Size*, *Age*, *Turnover*, *Risk*, and *Past Return*, defined in Table A1 of the Appendix. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	*Return*	*FF4 Alpha*	*Return*	*FF4 Alpha*	*Return*	*FF4 Alpha*
RAI Macro × *Post GPT*	2.196	2.086				
	(0.64)	(0.73)				
RAI Macro	0.727	-1.370*				
	(0.66)	(-1.71)				
RAI Firm Policy × *Post GPT*			4.840**	4.606***		
			(2.34)	(2.63)		
RAI Firm Policy			-0.367	-0.906		
			(-0.49)	(-1.36)		
RAI Firm Performance × *Post GPT*					2.331**	1.951**
					(2.37)	(2.30)
RAI Firm Performance					-0.294	-0.330
					(-0.92)	(-1.22)
Observations	9,974	9,974	9,974	9,974	9,974	9,974
R-squared	0.810	0.093	0.810	0.094	0.810	0.093
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 10. RAI and Hedge Fund Performance: ChatGPT Release

This table reports the relation between performance and reliance on AI information. *Return* is the portfolio holdings return. *CAPM Alpha* (*FF3 Alpha*/*FF4 Alpha*) is the portfolio holdings return, after adjusting for the market risk factor (Fama-French three factors/Fama-French-Carhart four factors). *RAI* is defined in equation (3), which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Post ChatGPT* is an indicator variable equal to one for performance in the first quarter of 2023 and onwards, and zero otherwise. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Turnover* is the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Past Return* is the lagged *Return*. The sample period is from the first quarter of 2016 to the second quarter of 2023 in Panel B. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

	(1)	(2)	(3)	(4)
Dep. Var.	*Return*	*CAPM Alpha*	*FF3 Alpha*	*FF4 Alpha*
RAI × *Post ChatGPT*	1.714**	2.983***	3.371***	2.596***
	(2.14)	(3.53)	(4.57)	(3.44)
RAI	-0.073	-0.168	-0.306	-0.289
	(-0.27)	(-0.61)	(-1.25)	(-1.15)
Size	-0.094***	-0.117***	-0.115***	-0.120***
	(-3.62)	(-3.26)	(-3.71)	(-3.91)
Age	0.013***	0.001	-0.001	-0.001
	(2.82)	(0.21)	(-0.20)	(-0.20)
Turnover	0.256	-0.482	-0.655**	-0.661**
	(1.05)	(-1.51)	(-2.26)	(-2.30)
Risk	9.201***	-6.734**	-14.063***	-13.695***
	(4.24)	(-2.51)	(-6.89)	(-6.27)
Past Return	0.100***	0.041***	-0.007	-0.009
	(5.88)	(3.16)	(-0.63)	(-0.84)
Observations	10,762	10,762	10,762	10,762
R-squared	0.798	0.107	0.094	0.089
Time FE	Yes	Yes	Yes	Yes

Table 11. Alternative RAI Measure

This table reports the relation between performance and reliance on AI information. *Return* is the portfolio holdings return. *CAPM Alpha* (*FF3 Alpha*/*FF4 Alpha*) is the portfolio holdings return, after adjusting for the market risk factor (Fama-French three factors/Fama-French-Carhart four factors). *Alt. RAI* is an alternative measure of *RAI* that quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. *Post GPT* is an indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise. *Size* is the natural logarithm of total holdings value. *Age* is the number of years since a hedge fund company's first 13F report. *Turnover* is the minimum of purchases and sales scaled by total holdings value. *Risk* is the standard deviation of quarterly portfolio returns in the past two years. *Past Return* is the lagged *Return*. The sample period is from the third quarter of 2022 to the second quarter of 2023 in Panel A and from the first quarter of 2016 to the second quarter of 2023 in Panel B. The *t*-statistics, in parentheses, are based on standard errors clustered by fund. *p <.1; **p <.05; ***p <.01.

Panel A: During Post-GPT period

	(1)	(2)	(3)	(4)
Dep. Var.	*Return*	*CAPM Alpha*	*FF3 Alpha*	*FF4 Alpha*
Alt RAI	0.073***	0.092***	0.094***	0.091***
	(2.73)	(3.74)	(3.93)	(3.92)
Size	-0.022	0.096	-0.055	-0.062
	(-0.25)	(1.02)	(-0.70)	(-0.80)
Age	-0.018	-0.040**	-0.025	-0.021
	(-1.04)	(-2.39)	(-1.56)	(-1.31)
Turnover	0.628	0.110	0.469	0.316
	(0.60)	(0.10)	(0.45)	(0.31)
Risk	5.458	4.821	1.008	1.668
	(1.09)	(1.16)	(0.25)	(0.39)
Past Return	-0.192***	-0.228***	-0.032	-0.037
	(-4.94)	(-6.35)	(-0.87)	(-0.98)
Observations	1,001	1,001	1,001	1,001
R-squared	0.452	0.115	0.038	0.041
Time FE	Yes	Yes	Yes	Yes

Panel B: DiD

Dep. Var.	(1) *Return*	(2) *CAPM Alpha*	(3) *FF3 Alpha*	(4) *FF4 Alpha*
Alt RAI × *Post GPT*	0.069**	0.089***	0.111***	0.107***
	(2.40)	(3.10)	(4.41)	(4.32)
Alt RAI	-0.012	-0.006	-0.013	-0.012
	(-1.14)	(-0.61)	(-1.61)	(-1.43)
Size	-0.094***	-0.117***	-0.111***	-0.116***
	(-3.72)	(-3.37)	(-3.73)	(-3.91)
Age	0.012***	0.001	-0.001	-0.001
	(2.76)	(0.16)	(-0.23)	(-0.20)
Turnover	0.230	-0.511*	-0.663**	-0.654**
	(1.00)	(-1.65)	(-2.33)	(-2.31)
Risk	9.258***	-6.782***	-14.291***	-13.974***
	(4.46)	(-2.66)	(-7.45)	(-6.80)
Past Return	0.100***	0.041***	-0.006	-0.009
	(5.88)	(3.18)	(-0.61)	(-0.84)
Observations	10,762	10,762	10,762	10,762
R-squared	0.798	0.107	0.094	0.090
Time FE	Yes	Yes	Yes	Yes

Appendix

Table A1. Definitions of Variables

Variable	Definition
Age	The number of years since a hedge fund company's first 13F report.
AI Hedge Fund	An indicator variable equal to one if a hedge fund has AI-skilled workers and zero otherwise.
Alt. RAI	An alternative reliance measure of *RAI*, defined as the $R^2_{AI,i,t} - R^2_{fundamental,i,t}$ (i.e., *RAI*) scaled by $R^2_{fundamental,i,t}$. See definition of *RAI* for details.
CAPM Alpha	At the end of quarter t, we use the monthly stock returns in the past 36 months to estimate the beta on the risk factor and calculate risk-adjusted return in quarter $t+1$ as the difference between realized stock return minus stock return estimated with beta. *CAPM Alpha* is the weighted average cumulative monthly abnormal return across all holdings, where the weight is the value of stock j held by i at the end of quarter t divided by the total value of all stocks held by i at the end of quarter t.
Early Adopter	An indicator variable equal to one if a hedge fund company has an insignificant F-statistic at the end of 2021 but a significant F-statistic at 1% level in 2022. F-statistic is defined as $F_{i,t} = ((RSS_{fundamental,i,t} - RSS_{AI,i,t})/p)/(RSS_{AI,i,t}/(n-k))$ where $RSS_{fundamental,i,t}$ is the residual sum of squares of the model with firm fundamentals only while $RSS_{AI,i,t}$ is the residual sum of squares of the full model after adding the fundamental information generated by ChatGPT. p is the number of predictors added to the full model and equals 14 since we have 14 ChatGPT scores. n is the number of observations and equals the number of trades in a given fund quarter. k is the number of coefficients (including the intercept) in the full model and equals 20 since we have five variables about firm fundamentals, 14 ChatGPT scores, and an intercept.
FF3 Alpha	The weighted average risk-adjusted returns using the Fama-French three-factor model. The construction is analogous to *CAPM Alpha*.
FF4 Alpha	The weighted average risk-adjusted returns using the Fama-French-Carhart four-factor model. The construction is analogous to *CAPM Alpha*.
Hedge Fund	An indicator variable if an investment company is a hedge fund company and zero otherwise.
Outage	An indicator variable equal to one during the quarters with higher than median ChatGPT outage occurrences, and zero otherwise.
Past Return	The one-quarter-lagged *Return*.
Post ChatGPT	An indicator variable equal to one for performance in the first quarter of 2023 and onwards, and zero otherwise.
Post GPT	An indicator variable equal to one for performance in the third quarter of 2022 and onwards, and zero otherwise.

(continued)

Variable	Definition
RAI	Reliance on generative AI information, which quantifies the degree to which changes in portfolio holdings are influenced by fundamental information generated by ChatGPT. We estimate *RAI* using a two-step procedure. In the first step, at the end of each quarter t and for each investment company i, we run the following two regression models across the investment company's stock j trades in quarter t: $HoldingChange_{i,j,t} = \gamma_{i,t} \cdot X_{j,t-1} + \varepsilon_{i,j,t}$ and $HoldingChange_{i,j,t} = \Sigma_{j=1}^{J} \beta_{i,t} \cdot GPT\ Score_{j,t-1} + \gamma_{i,t} \cdot X_{j,t-1} + \varepsilon_{i,j,t}$ where $X_{j,t-1}$ is a host of financial variables about firm fundamentals in quarter $t-1$, including market capitalization, book-to-market, return on assets, stock return, and change in the analyst recommendation consensus. *GPT Score* includes 14 signals generated by ChatGPT, covering firms' expectations about macroeconomic, industry, and firm-specific performance and policy outcomes. A full list of signals is in Table A2 in the Appendix. We define the R^2 from the first equation as $R^2_{fundamental,i,t}$, and the R^2 from the second equation as $R^2_{AI,i,t}$, and $RAI_{i,t} = R^2_{AI,i,t} - R^2_{fundamental,i,t}$.
RAI Firm Performance	A decomposed *RAI* measure and the construction is analogous to *RAI*, except that *GPT Score* only includes signals about a firm's earnings, revenue, financial prospects, and product market.
RAI Firm Policy	A decomposed *RAI* measure and the construction is analogous to *RAI*, except that *GPT Score* only includes signals about a firm's wages, employment, capital expenditure, and cost of capital.
RAI Macro	A decomposed *RAI* measure and the construction is analogous to *RAI*, except that *GPT Score* only includes signals about the global economy, the US economy, and a firm's industry.
Return	The weighted average cumulative monthly return across all holdings in quarter $t+1$, where the weight is the value of stock j held by i at the end of quarter t divided by the total value of all stocks held by i at the end of quarter t.
Risk	The standard deviation of quarterly portfolio returns in the past two years.
Size	The natural logarithm of total holdings value.
Turnover	The minimum of purchases and sales over average total holdings values of the current quarter and the previous quarter, following Carhart (1997).

Table A2. List of Questions to Generate AI Information

This table reports the list of questions used to query ChatGPT and generate forward-looking information/signal based on firms' earnings conference call transcripts. These questions are based on Jha, Qian, Weber, and Yang (2023, 2024).

	Over the next quarter, how does the firm anticipate a change in:
No.	Topic
1	optimism about the US economy?
2	optimism about the global economy?
3	optimism about the financial prospects of their firm?
4	optimism about the financial prospects of its industry?
5	its earnings?
6	its revenue?
7	its wages and salaries expenses?
8	demand for its products or services?
9	production quantity of its products?
10	prices for its products or services?
11	prices for its inputs or commodities?
12	its cost of capital or hurdle rate?
13	its capital expenditure?
14	its employment?