Do Information Processing Frictions Drive the Local Bias?

Angie Andrikogiannopoulou* Filippos Papakonstantinou†
King's College London King's College London

Hao Song‡
King's College London

Abstract

We provide causal evidence that the cost of processing public information is a quantitatively important driver of investors' local bias. Exploiting the public release of capable generative AI as a shock that lowered the cost of interpreting and synthesising public disclosures, we examine its effect on mutual funds' allocations to geographically proximate and distant firms. Using a difference-in-differences design, we find that funds' relative overweighting of local firms declines by about one-third after this shock. The effect is concentrated among firms with lengthy and complex disclosures and among discretionary and non-team-managed funds, where ex ante processing constraints are more likely to bind. Consistent with generative AI relaxing the constraints associated with analysing geographically distant firms, we further find that the stock-picking ability of distant portfolios improves relative to that of local portfolios.

*King's Business School, Department of Banking & Finance, Bush House, 30 Aldwych, London WC2B 4BG, UK, e-mail: aandriko@kcl.ac.uk.

†King's Business School, Department of Banking & Finance, Bush House, 30 Aldwych, London WC2B 4BG, UK, e-mail: fpapakon@kcl.ac.uk.

‡King's Business School, Department of Banking & Finance, Bush House, 30 Aldwych, London WC2B 4BG, UK, e-mail: hao.song@kcl.ac.uk.

1 Introduction

A large body of literature documents that institutional investors systematically overweight geographically proximate firms in their portfolios, commonly referred to as local bias (Coval and Moskowitz, 1999; Coval and Moskowitz, 2001; Baik, Kang and Kim, 2010; Bernile, Kumar and Sulaeman, 2015)[1]. This geographic tilt represents one of the persistent distortions in portfolio allocation. Standard portfolio theory implies that investors should hold diversified portfolios based on firms' fundamentals and risk characteristics, rather than on where firms are headquartered relative to the investor. By systematically tilting toward nearby firms, investors may forgo the benefits of diversification and increase their exposure to correlated local economic shocks, such as regional employment or industry-cluster shocks that affect both the investor's local economy and nearby firms (Korniotis and Kumar, 2013). Yet the underlying cause of local bias remains debated. A leading interpretation is that proximity confers an information advantage, along three distinct margins: local investors may be better able to acquire, access, or process firm-specific information[2]. Despite the prominence of this explanation, direct causal evidence linking information advantage to local bias has remained limited. The existing evidence has focused on the first two margins—the acquisition of private or soft information (e.g., through site visits, see Cicero et al., 2024) or the supply and dissemination of public information (e.g., through local news, see Kang and Nam, 2025)—leaving the role of information processing frictions largely unresolved.

Our paper focuses on this third, under-explored margin. Theoretical work by Van Nieuwerburgh and Veldkamp (2009; 2010) and Kacperczyk, Van Nieuwerburgh and Veldkamp (2016) argues that investors have limited capacity to process information. Subject to this constraint, investors rationally allocate attention toward assets that are less costly to learn about, and local firms tend

[1]Local bias is also well documented among retail investors. Individual investors disproportionately hold and trade stocks headquartered near where they live; see, for example, French and Poterba (1991), Ivkovic and Weisbenner (2005), Bodnaruk (2009), Seasholes and Zhu (2010).

[2]Conceptually, information frictions can operate along at least three distinct margins. First, investors may differ in their ability to acquire private or soft information, such as through site visits or personal networks. Second, geography may shape the supply and dissemination of public information, for example through local media coverage. Third, even when information is publicly available to all investors, they may differ in the cost of processing that information—that is, in their ability to extract, contextualise, synthesise, and interpret firm disclosures or other publicly available information.

to satisfy this criterion. Nearby investors typically have greater familiarity with a local firm's operational and economic environment, including its customers, suppliers, competitors, labour market, and regional regulatory conditions[3]. This contextual knowledge can make the same public disclosure easier to interpret for local investors than for distant investors[4]. Under this information-processing-frictions view, local bias should attenuate when a shock lowers the cost of processing information about distant firms, even if geographic proximity itself remains unchanged. Testing this prediction, however, requires a plausibly exogenous change in information-processing costs.

Anecdotal evidence suggests that generative artificial intelligence (GenAI) began entering professional investment workflows soon after ChatGPT's release. In early 2023, Morgan Stanley partnered with OpenAI and rolled out a GenAI assistant to help financial advisers access the firm's research and intellectual capital, while Bloomberg introduced BloombergGPT, a finance-specific large language model trained on financial data[5]. Moreover, emerging academic evidence points in the same direction. A growing literature shows that investors are rapidly adopting GenAI in their investment process (Cheng, Lin and Zhao, 2025; Sheng et al., 2026). Specifically, Sheng et al. (2026) document a sharp increase in asset managers' reliance on GenAI in making investment decisions after ChatGPT's release. At the same time, recent studies show that the use of GenAI can materially reduce the cost of processing public textual information (Bilinski, 2024; Lopez-Lira and Tang, 2026), for example by reducing the cost of extracting, contextualising, and synthesising textual information from 10-K filings, analyst reports, and news. In line with this view, survey evidence in Sheng et al. (2026) shows that the vast majority of GenAI-adopting funds use these tools specifically to process and analyse financial text.

Motivated by this evidence, we exploit the emergence of *capable* GenAI, marked by the public

[3]In this context, familiarity refers to awareness and contextual knowledge about a local firm's economic environment, as opposed to a taste-based preference for nearby firms. This notion aligns with Kuchler et al. (2022), who emphasise social proximity as increasing investors' awareness of firms, and differs from the familiarity-preference channel in Pool, Stoffman and Yonker (2012), where investors favour home-state firms because they feel more comfortable with them, even absent superior information.

[4]In addition, geographic proximity may facilitate informal information flows through local media, social networks, site visits, or direct contact with management.

[5]See: https://www.morganstanley.com/press-releases/key-milestone-in-innovation-journey-with-openai; https://arxiv.org/abs/2303.17564.

release of ChatGPT[6], as a technological shock to information-processing costs and provide causal evidence on the role of information-processing frictions in driving local bias. Importantly, this technological change does not directly alter the production and availability of public or private information. It therefore allows us to isolate information-processing advantages from alternative sources of local advantage. Our central hypothesis follows directly: if local bias partly reflects information-processing costs that favour local investors, then the emergence of GenAI should disproportionately lower the cost of evaluating distant firms, where information-processing frictions are initially greater[7], and lead to a decline in local overweighting.

We examine this prediction using detailed, frequently updated holdings of U.S. active equity mutual funds. These data are well suited to our question because they allow us to observe how professional investors allocate capital between local and distant stocks over time, and how this allocation shifts around ChatGPT's release. The sample period spans from the second quarter of 2021 to the second quarter of 2024, capturing a balanced window around the shock. Variation in fund–firm distance, which is largely fixed during the sample period, provides a natural source of differential exposure to potential reductions in information-processing costs, and this enables a difference-in-differences (DiD) design to test whether these reductions weaken local overweighting. Following the literature (e.g., Alok, Kumar and Wermers, 2020; Lee, 2026), we classify a firm as local to a fund if the firm's headquarters is located within 100 miles of the fund adviser's address[8].

We begin our analysis by examining changes in the portfolio weights that mutual funds assign to local and distant firms around ChatGPT's release. Our analysis is conducted at the

[6]Large language models existed before ChatGPT, including BERT, GPT-2, GPT-3, and T5. However, these earlier models were either designed primarily for task-specific fine-tuning or required substantial technical expertise to use effectively, limiting their practical relevance for ordinary investment workflows. ChatGPT marked the emergence of capable GenAI because it made large language models broadly accessible through a conversational interface and improved their ability to follow instructions through instruction tuning and reinforcement learning from human feedback. This made GenAI practically useful for general-purpose text-processing tasks, such as extracting, summarising, translating, and synthesising firm-level information.

[7]The marginal benefit of GenAI is likely to be greater for distant firms. Local investors already possess contextual knowledge about nearby firms. This familiarity reduces the marginal value of additional tools for processing public information. By contrast, distant firms are initially more costly to evaluate because investors lack this local context. GenAI can therefore have a larger marginal effect on distant firms by helping investors contextualise and interpret information that would otherwise be more costly to process.

[8]Our results are robust to alternative cutoffs, including the 150-mile threshold and 250-mile threshold used by Ivkovic and Weisbenner (2005).

fund–firm–quarter level. To the extent that GenAI disproportionately lowers the cost of processing information about distant firms, we expect the local–distant portfolio-weight gap to narrow after ChatGPT becomes publicly available. Consistent with this prediction, we find that funds overweight geographically proximate firms before ChatGPT's release, but this local overweighting declines by more than a third thereafter, controlling for fund, firm, and year-quarter fixed effects. The result is robust to controlling for a comprehensive set of time-varying fund and firm characteristics and also survives increasingly demanding fixed-effect structures, including fund-by-industry fixed effects and fund-by-firm fixed effects. The stability of the estimates across these specifications alleviates concerns that the decline in local overweighting is driven by observable fund or firm characteristics, persistent industry allocation patterns, or stable fund–firm matching.

However, we recognise that the decline in local overweighting could mechanically reflect a post-ChatGPT shift toward, or stronger performance of, technology stocks rather than a reduced reliance on geographic proximity. As many funds in our sample are located in New York and Boston, whereas technology firms are disproportionately concentrated on the West Coast, either channel would mechanically increase the portfolio weights of distant stocks and thereby reduce the measured local–distant weight gap. We address this concern in several ways. First, we re-estimate the baseline specification after separately excluding funds headquartered in New York or Boston and firms in the technology sector, and obtain similar results. Second, we control for contemporaneous stock returns to account for mechanical valuation effects on portfolio weights; the estimated decline in local overweighting remains robust. Finally, we introduce fund-by-year-quarter and firm-by-year-quarter fixed effects, which absorb time-varying fund behaviour and firm-specific shocks in each quarter. The estimated effect remains negative and statistically significant, suggesting that our findings are not explained by a post-ChatGPT rotation toward particular industries or by the relative appreciation of distant stocks.

To further assess the validity of our empirical design and trace the evolution of local overweighting around ChatGPT's release, we estimate an event-study specification. The pre-event coefficients show no evidence of differential trends in local and distant portfolio weights before the shock, sup-

porting the parallel-trends assumption underlying our difference-in-differences design. Following the release, the decline in local overweighting emerges gradually rather than immediately: the effect is small and insignificant in the first post-event quarter, becomes negative from the second quarter onward, and strengthens over subsequent quarters. This pattern is consistent with the gradual diffusion of GenAI into investment workflows and with improvements in model capability over time, including the releases of GPT-4 and GPT-4 Turbo. Overall, the dynamic evidence reinforces the baseline findings and suggests that the reduction in local bias intensifies as GenAI becomes more capable and more widely incorporated into fund managers' investment processes.

The effect of GenAI emergence on portfolio adjustment may work through the extensive margin, by changing the set of stocks funds choose to hold, or through the intensive margin, by changing how funds allocate capital across their existing holdings—or through both. We therefore assess the contribution of each margin to the decline in local bias. We first find that, after ChatGPT's release, the probability of mutual funds holding distant stocks increases relative to that of holding local stocks. Decomposing this change, we find that funds are relatively more likely to initiate new positions in distant stocks, while there is no evidence that they become more likely to fully liquidate their existing local positions. Beyond these changes in holding decisions, we also find that, among continuing positions, funds reallocate capital away from local stocks and toward distant ones. Taken together, both margins contribute to the decline in local bias, with the extensive margin operating primarily through new entry into distant stocks rather than complete exit from existing local holdings. This pattern is consistent with the idea that GenAI lowers the information-processing barriers to evaluating distant firms, while leaving some value in previously accumulated local information advantages.

Next, we examine how the post-ChatGPT decline in local overweighting varies across firm and fund characteristics that proxy for the potential usefulness of GenAI, providing further evidence on the information-processing mechanism. If GenAI reduces local bias by lowering the cost of processing firm-level information, its effect should be stronger where such costs are initially higher and where investors face tighter information-processing constraints. We first consider firm-level disclosure complexity. Using firms' 10-K filings, we measure complexity by filing size and the

SMOG readability index. Consistent with the proposed mechanism, the decline in local overweighting is concentrated among firms with more complex disclosures, while the effect is small and statistically insignificant for firms with relatively simple disclosures.

We then turn to exploring heterogeneity across fund investment strategies. Building on Abis (2022), we classify funds as discretionary, where investment decisions rely heavily on managerial judgement, or quantitative, where investment decisions are purely model-driven. If GenAI complements human judgement by expanding managers' ability to process and interpret information, its effects should be more evident among discretionary funds. Consistent with this view, the decline in local overweighting is statistically significant among discretionary funds, but not among quantitative funds. Additionally, we examine managerial structure as a proxy for *ex ante* information-processing capacity. The decline in local overweighting is larger in magnitude and statistically significant among non-team-managed funds, whereas it is smaller and statistically insignificant among team-managed funds. Because team-managed funds may already possess greater processing capacity through specialisation and division of labour, this pattern is consistent with GenAI being particularly useful when managerial processing capacity is more limited. Taken together, these cross-sectional results corroborate our interpretation that the decline in local bias reflects the information-processing benefits of GenAI.

We next ask whether the post-ChatGPT reallocation toward distant firms is informative. If GenAI improves managers' ability to process firm-specific information about distant stocks, the additional capital allocated to them should flow disproportionately toward those identified as attractive investment opportunities. We thus divide each fund's portfolio into local and distant holdings and assess how managers' stock-picking ability, measured following Kacperczyk, Van Nieuwerburgh and Veldkamp (2016), changes across the two segments around ChatGPT's release. We find that, before the release, local holdings exhibit significantly higher stock-picking ability than distant holdings, consistent with a pre-existing informational advantage associated with geographic proximity. After the release, this local stock-picking advantage declines significantly. In other words, the post-ChatGPT improvement in the informativeness of distant holdings exceeds that of local

holdings. Further, we examine whether the relative improvement in the informativeness of distant allocations translates into realised portfolio performance. We show that, after ChatGPT's release, the performance of distant holdings improves relative to that of local holdings. The local-minus-distant return gap declines for both excess returns and DGTW-adjusted returns, by approximately 1.6 percentage points and 0.3 percentage points, respectively. Overall, the performance evidence complements our portfolio-weight results and is consistent with our argument that GenAI reduces the information-processing disadvantage associated with geographic distance, thereby improving the quality of investment decisions in distant stocks.

Related Literature

Our paper builds on the extensive literature documenting the drivers of local investment bias. Coval and Moskowitz (1999) propose—and Van Nieuwerburgh and Veldkamp (2010) formalise—the interpretation that geographic proximity lowers information frictions, and so local bias reflects information asymmetries rather than purely behavioral preferences. Consistent with this view, a number of studies find that investors earn abnormal returns on local investments, suggesting that local bias can reflect genuine informational advantages. For example, Ivkovic and Weisbenner (2005), Massa and Simonov (2006), and Ben-David, Birru and Rossi (2019) find such evidence for retail investors, while Coval and Moskowitz (2001), Hau (2001), Teo (2009), and Sialm, Sun and Zheng (2020) find similar evidence for institutional investors. At the same time, another body of work finds no abnormal performance on local holdings, especially among less sophisticated investors (Grinblatt and Keloharju, 2001; Seasholes and Zhu, 2010; Pool, Stoffman and Yonker, 2012), and proposes non-information explanations for local bias, including familiarity, loyalty, patriotism, and trust (Huberman, 2001; Cohen, 2009; Morse and Shive, 2011; Shao and Wang, 2021). Together, this literature suggests that local bias may not reflect a single mechanism, but might arise from different informational and non-informational channels.

We contribute to this debate by providing new causal evidence that not only supports the information-based view of local bias but also isolates a specific informational margin that has

remained empirically elusive. Conceptually, the literature points to three distinct channels through which geography may shape portfolio choice: (i) the supply of public information, (ii) the acquisition of private information, and (iii) the cost of processing public information. Existing causal or quasi-causal evidence primarily speaks to the first two channels. For example, Kang and Nam (2025) exploit staggered closures of local newspapers to show that local informational advantages widen when the supply of public information declines, while Da et al. (2021) and Bai and Massa (2026) use shocks to travel easiness and face-to-face interactions—such as airline route introductions and COVID-19 lockdowns—to demonstrate that local advantages arise from active information acquisition through site visits.[9] We complement this literature by exploiting a plausibly exogenous large-scale technological shock—the emergence of GenAI—to study the third channel: the cost of information processing. By isolating this margin, our analysis sheds new light on the informational foundations of local bias. Our results show that, reductions in public-information processing costs attenuate local bias, implying that geographic proximity can serve as a substitute for the costly interpretation of complex disclosures[10].

Moreover, our paper relates to a growing literature on the role of information technology in financial markets. Prior work shows that technological advances such as improved internet access and the electronic dissemination of financial information—including through the EDGAR system and social media platforms—affect investor participation, information acquisition, and

[9]There is also a broader literature providing causal evidence on the role of information frictions in financial markets. Some studies examine shocks to the supply and dissemination of public information—such as changes to disclosure regulations (e.g., Bushee and Leuz, 2005) or median coverage (e.g., Peress, 2014)—and show that investor trading and price efficiency respond to these changes. Other work focuses on shocks to private information acquisition and monitoring costs—such as reduction in travel times (Bernstein, Grioud and Townsend, 2016) and weather-induced shocks to IPO roadshow interactions (Gustafson et al., 2025)—and show that these changes affect investor involvement and firm outcomes.

[10]The distinct informational channels discussed above differ in their economic scope and in their ability to explain local bias. First, the information-supply channel emphasizes the role of geography in shaping variation in information availability and, consequently, in investment behavior. However, in modern financial markets—characterized by extensive disclosure and rich information environments—an important challenge faced by investors is often not access to information, but the capacity to process it effectively (Kacperczyk, Van Nieuwerburgh and Veldkamp 2016; Blankespoor et al. 2018). Second, the information-acquisition channel, which operates through in-person or social interactions, is inherently rivalrous and costly, and therefore difficult to scale across investors and firms. In contrast, differences in the ability to process publicly available information operate at a broader margin, affecting all investors and applying broadly across firms. By isolating the information-processing margin, our paper evidences a key determinant of local bias as a pervasive and persistent portfolio distortion observed across investors, asset classes, and time periods, rather than as localized alpha opportunities available to a small subset of investors.

market efficiency (e.g., Barber and Odean, 2002; Blankespoor, Miller and White, 2014; Gao and Huang, 2020; Xu, Xuan and Zheng, 2021; Chang, Ljungqvist and Tseng, 2023). Related studies of the availability of machine-readable quantitative disclosures, such as XBRL, and the use of textual-analysis techniques to process unstructured disclosures show that reductions in information-extraction and processing costs facilitate the incorporation of information into prices (Loughran and MdConald, 2011; Dong et al., 2016; Blankespoor, 2019; Cohen, Malloy and Nguyen, 2020; Zeng and Zhou, 2026). While conceptually close, most prior technologies either affected investors systematically or coincided with changes in disclosure regimes themselves. In contrast, GenAI represents a qualitatively different innovation: it substantially lowers information-processing costs by directly assisting investors in interpreting complex textual disclosures, with marginal benefits that vary across investors and firms, while leaving disclosure requirements unchanged. By reducing the time and cognitive effort required to process information, GenAI effectively expands investors' information capacity, enabling them to evaluate firms that would otherwise be relatively costly to analyse. Given these features, we exploit the emergence of GenAI as a novel empirical setting to study how reductions in public-information processing frictions reshape portfolio allocation.

Finally, our paper speaks to the nascent literature that examines how GenAI affects financial-market participants and market outcomes. Recent studies show that GenAI enhances analysts' forecast performance, facilitates investor trading, and improves professional fund managers' analysis of firm-specific information and investment performance (e.g. Bertomeu et al. 2025, Cheng, Lin and Zhao, 2025; Sheng et al., 2026). Complementary evidence demonstrates that large language models can extract economically meaningful signals from financial news and use them to predict subsequent stock returns (Lopez-Lira and Tang, 2026). We extend this literature by examining how post-GenAI information environment affects institutional investors' actual portfolio choices. We show that the emergence of GenAI constitutes a disruptive innovation that reshapes the informational environment faced by investors and attenuates a long-standing portfolio distortion—local bias—by reducing the relative information-processing disadvantage associated with distant firms. Importantly, this reallocation is informative: increased investment in distant firms is accompanied

by improved stock-picking skill. Our findings therefore show that GenAI affects not only where investors allocate capital, but also the informational content of that reallocation.

2 Institutional background and data

2.1 Background and hypothesis

The recent emergence of GenAI represents a disruptive development in information-processing technology. Advances in large language models substantially improved machines' ability to understand natural-language prompts, extract information from unstructured text, and generate coherent responses across a wide range of tasks. Earlier language models, such as GPT-2 and early versions of GPT-3, had already demonstrated some capabilities in text generation, summarisation, and question answering. However, their performance on complex tasks was more limited. Moreover, their practical use often required greater technical expertise or specialised interfaces, limiting their accessibility to a broader set of users.

ChatGPT's debut on November 30, 2022 marked an important turning point. Built on the GPT-3.5 model family and fine-tuned to follow user instructions and engage in conversational interaction, ChatGPT combined substantially improved usability with stronger capabilities for performing complex language-based tasks. Through a simple chat interface, users could provide detailed natural-language instructions, refine questions iteratively, and apply the model to tasks such as summarising lengthy documents, extracting relevant information, comparing and contextualising arguments, and synthesising information across different pieces of text. These capabilities are particularly relevant for information-intensive professional activities in which the principal constraint is not necessarily access to information, but the time and effort required to process it. This combination of ease of use and practical capability led to exceptionally rapid adoption: ChatGPT reportedly reached one million users within five days of its release and approximately 100 million users within two months[11].

Following ChatGPT's initial release, the technology continued to evolve rapidly. OpenAI in-

[11] Source: https://openai.com/global-affairs/new-economic-analysis/.

troduced ChatGPT Plus in February 2023 and made ChatGPT available through an API in March 2023. GPT-4 was released on March 14, 2023, offering materially stronger performance on a broad range of language and reasoning tasks. This was followed by GPT-4 Turbo in November 2023, which expanded the amount of information that could be processed within a single interaction while reducing usage costs. GPT-4o, released in May 2024, further improved speed and multimodal capabilities. These successive developments progressively increased the capability, accessibility, and scalability of GenAI, making it increasingly suitable for professional information-processing tasks. Figure 1 summarises the major milestones in the evolution of ChatGPT and related OpenAI models during our sample period.

This institutional development is particularly important for the asset-management industry, where investment research relies not only on processing structured financial data but also on interpreting large volumes of unstructured textual information, including corporate disclosures, analyst reports, and financial news. Given the central role of information processing in asset management, the emergence of GenAI—which can lower the cost of extracting, interpreting, and synthesizing information across numerical and textual sources—provides a natural setting to examine an information-processing-frictions explanation for local investment bias, as illustrated below.

Van Nieuwerburgh and Veldkamp (2009; 2010) and Kacperczyk, Van Nieuwerburgh and Veldkamp (2016) show that investors have limited information-processing capacity and therefore cannot evaluate all assets with equal depth. Subject to this constraint, investors allocate scarce attention toward assets for which they possess an initial informational advantage or for which information is less costly to process. Geographic proximity can generate such an advantage because investors located near a firm are likely to possess greater familiarity with its economic and institutional environment, including its customers, suppliers, competitors, labour market, and regional conditions. This contextual knowledge can make the same firm-level information easier to interpret for a local investor than for a geographically distant investor. Local investment bias, from this perspective, does not necessarily reflect a preference for proximity per se, but instead arises from the endogenous allocation of limited attention toward firms with lower informational costs. Although

this framework provides a theoretical explanation for local investment bias, it has not been directly examined empirically.

The emergence of GenAI directly affects this information-processing mechanism by reducing the time and effort required to extract, interpret, contextualise, and synthesise firm-level information. In doing so, it effectively expands investors' information-processing capacity. Because information-processing frictions are initially greater when evaluating distant firms, there is greater scope for GenAI to reduce these frictions. By helping investors interpret complex disclosures, connect information across sources, and provide context that would otherwise be costly to acquire and process, GenAI can narrow the information-processing gap between local and distant investors. As a result, fund managers should become less dependent on the information-processing advantage associated with geographic proximity and more willing to allocate capital to distant firms. The release of ChatGPT thus provides a discrete shock to the information-processing environment, while subsequent improvements in GenAI capabilities may progressively strengthen its economic effects.

We focus on the public release of ChatGPT rather than the introduction of other GenAI technologies for several reasons. First, ChatGPT represents a particularly salient and broadly accessible introduction of highly capable conversational AI to the general public. As discussed above, its adoption was exceptionally rapid. Second, the timing of its release preceded the broad public availability of several prominent alternative conversational AI systems[12]. Using the release of these later systems would make it difficult to distinguish their effects from the information-processing changes already initiated by ChatGPT. Importantly, our interpretation does not require investors to use ChatGPT specifically. Rather, we view its release as marking the beginning of a broader shift in the information-processing environment, after which competing models and specialised applications proliferated. Accordingly, our treatment captures the effect of the emergence and diffusion of broadly accessible GenAI rather than the adoption of a particular GenAI product.

[12]Meta released the first LLaMA models in February 2023 primarily for the research community, Google began opening public access to Bard in the United States and United Kingdom on 21 March 2023, and Anthropic introduced Claude more broadly in July 2023 after an earlier closed-alpha period.

2.2 Mutual fund and firm data

We draw on multiple data sources to construct our dataset, as described below.

We obtain fund characteristics from Center for Research in Security Prices (CRSP) Survivorship-Bias-Free U.S. Mutual Fund Database, including total net assets (TNA), net returns, portfolio turnover, expense ratios, Lipper classification codes, CRSP objective codes and other fund attributes. Our sample consists of actively managed U.S. domestic equity funds from the second quarter of 2021 to the second quarter of 2024. We begin with funds falling within Lipper classification codes corresponding to Large-Cap Core, Large-Cap Growth, Large-Cap Value, Mid-Cap Core, Mid-Cap Growth, Mid-Cap Value, Multi-Cap Core, Multi-Cap Growth, Multi-Cap Value, Small-Cap Core, Small-Cap Growth, Small-Cap Value, and Equity Income. We further apply filters based on CRSP objective codes and fund flags, together with keyword searches of fund names, to exclude fixed-income, money market, balanced, index, exchange-traded, and variable annuity funds[13]. Following prior studies, we exclude funds with fewer than 20 or more than 500 holdings, as well as funds with total net assets below $5 million.

We also collect fund adviser names from CRSP and identify their addresses using the SEC Form ADV Part 1 data files, which provide information on investment advisers and their business operations[14]. We match advisers across the two data sources using a name-matching algorithm. We focus on the location of the fund adviser rather than that of the management company because the adviser is the entity responsible for managing the fund's portfolio and making investment decisions[15]. As such, adviser location more closely captures where the investment decisions are made—and thus the relevant source of local informational advantages—than the fund's management company address. Nevertheless, when adviser information is missing or an adviser's address

[13]We exclude exchange-traded funds and variable annuity funds using the *et_fund* and *vau_fund* variables in CRSP. To exclude index funds, we use the *index_fund_flag* variable, search fund names for the terms "INDEX", "S&P", "DOW 30", and "NASDAQ", and remove funds managed by fund families known to follow index-like strategies, including Dimensional Fund Advisors, Direxion, Potomac, ProFunds, and Rydex.

[14]See: https://www.sec.gov/data-research/sec-markets-data/information-about-registered-investment-advisers-exempt-reporting-advisers.

[15]The management company primarily provides organisational and administrative support to the fund complex, while portfolio investment decisions are typically made by the fund adviser.

cannot be identified, we instead use the fund's management company address provided by CRSP.

We next turn to data on funds' portfolio holdings. CRSP provides detailed information on funds' portfolio holdings, including the identity of each security held, the number of shares held, the market value of each position, and each holding's percentage of TNA. We supplement these data with portfolio firms' fundamental information, such as firm size, book-to-market ratio, profitability, and financial leverage, from Compustat; market information, such as stock prices and returns, from CRSP; and analyst coverage from I/B/E/S. Firms' historical headquarters addresses are sourced from the SEC Analytics Suite provided by WRDS, a database constructed from information in firms' SEC filings that WRDS systematically parses and standardises to track firm information over time. Since our analysis focuses on the domestic portion of funds' portfolios, we exclude holdings in firms headquartered outside the United States. As in Lee (2026), we also exclude funds and firms located in Alaska, Hawaii, Puerto Rico, and the U.S. Virgin Islands to mitigate the influence of extreme geographic distances. We additionally omit firms with an IPO or delisting during the sample period to ensure that the set of investable firms remains consistently observable throughout the analysis and that changes in portfolio holdings are not mechanically driven by firms entering or exiting the public equity market.

We geocode the fund adviser's address and the firm's headquarters address into latitude and longitude coordinates using U.S. Census Bureau's Geocoding Services API[16]. We then compute the great-circle distance between each fund–firm pair using the Haversine formula, which accounts for the curvature of the Earth and is given by:

$$d_{i,j} = 2r \arcsin\left(\sqrt{\sin^2\left(\frac{\phi_i - \phi_j}{2}\right) + \cos(\phi_i)\cos(\phi_j)\sin^2\left(\frac{\lambda_i - \lambda_j}{2}\right)}\right), \tag{1}$$

where $d_{i,j}$ is the distance between fund i's advisor and firm j's headquarters, r is the Earth's radius ($3,959$ miles), ϕ_i and ϕ_j denote latitude, and λ_i and λ_j denote longitude, with all coordinates expressed in radians[17]. We classify a portfolio firm as local to a fund if its headquarters is within

[16]See: https://geocoding.geo.census.gov/geocoder/locations/onelineaddress.

[17]Our results are robust to alternatively calculating geographic distance using the spherical law of cosines, as in Coval and Moskowitz (1999).

100 miles of the fund adviser, and as distant otherwise. To maintain a consistent local–distant classification over time, we drop fund–firm pairs whose geographic distance changes over the sample period in a way that alters their classification between local and distant. The final sample consists of $1,482$ distinct funds and $3,251$ unique stock holdings.

In addition, we construct several other fund-level variables by aggregating share-class data[18]. Although different share classes of the same fund hold identical portfolios, they may differ in TNA, fees, and net returns. Specifically, fund TNA is defined as the sum of TNA across all share classes, while net returns and expense ratios are calculated as TNA-weighted averages across classes, using beginning-of-quarter TNA as the weight. We calculate quarterly fund flows as: $Flow_{i,t} = \frac{TNA_{i,t} - TNA_{i,t-1}(1+r_{i,t})}{TNA_{i,t-1}}$, where $TNA_{i.t}$ is fund i's total net assets at the end of quarter t and $r_{i,t}$ is its net return in quarter t. We further classify each fund as targeted to institutional or to retail investors, based on CRSP's institutional share-class indicator for the share class with the largest TNA.

Summary statistics

Table 1 reports the average quarterly composition of the sample of funds across Lipper investment categories. The sample contains an average of $1,269$ funds per quarter, representing approximately \$3.12 trillion in aggregate TNA. In terms of the number of funds, however, Small-Cap Core is the most represented category, with an average of 171 funds per quarter, whereas Mid-Cap Value has the fewest, with 32 funds. Average fund size also varies considerably across categories. Large-Cap Core and Large-Cap Growth have the largest average fund TNA, at approximately \$5.47 billion and \$5.45 billion, respectively, while Small-Cap Value has the smallest, at approximately \$532 million.

The unit of observation in our analysis is a fund–firm–quarter. Our fund–firm sample includes firms that the fund held at any point during the sample period, as well as firms held by at least 25% of peer funds in the same Lipper investment style prior to the ChatGPT launch[19]. Consistent with

[18]We match CRSP share class identifiers to the corresponding fund identifiers in SEC filings (CIK–Series ID pairs) using the CRSP_CIK_MAP file from Wharton Research Data Services.

[19]The peer-based component of the investment universe is constructed using peer funds' pre-ChatGPT holdings so that it reflects a predetermined set of feasible investment opportunities and is not mechanically affected by changes in peer holdings following the shock. Our results are also robust to defining the peer-based investment universe more conservatively using a 50% peer-holding threshold.

the approach in Cremers and Petajisto (2009) and Lee (2026), we assign a portfolio weight of zero in quarters when the fund held no position in a given stock. Table 2 reports descriptive statistics for the main variables used in our analysis. Panel A shows that the average portfolio weight is 0.454%. The median portfolio weight is zero, compared with 1.490% at the 90th percentile, highlighting the considerable variation in portfolio allocations across firms within a fund's relevant investment universe. Panel B reports summary statistics for fund characteristics. The average fund flow and quarterly return are approximately 7% and 3%, respectively, while the average expense ratio is approximately 1%. Panel C presents the corresponding statistics for firm characteristics.

To illustrate the geographic distribution of funds and firms in our sample, we compute, for each quarter, the share of fund advisors and the share of firms located in each state, and report the time-series average of these shares in Figure 3. Fund advisors' locations are concentrated in a small number of states. Massachusetts hosts the largest share of funds, at approximately 17.8%, followed by New York at 15.7% and California at 9.3%; these three states alone account for 42.9% of all funds. This concentration reflects the clustering of asset management in established financial centers, where funds benefit from access to deep pools of investment talent and specialized financial services. Besides, funds are relatively well represented in Illinois, Connecticut, Texas, Colorado, Missouri, Ohio, and Pennsylvania.

Firm locations are somewhat more geographically dispersed, although several states account for a substantial fraction of firms in the sample. California has the largest share, containing approximately 17.1% of firms, followed by New York and Texas, which account for 10.4% and 10.0%, respectively. Massachusetts, Pennsylvania, Florida, and Illinois also contain relatively large proportions of firms. The concentration of firms in these states is broadly consistent with the uneven geographic distribution of economic and corporate activity across the United States. Importantly, the geographic distribution of firms does not simply mirror that of funds. The partial overlap between fund and firm locations creates meaningful variation in geographic proximity across fund–firm pairs, which is central to our analysis of local investment bias.

3 Does the emergence of GenAI reduce mutual funds' local overweighting?

In this section, we provide causal evidence on whether the emergence of GenAI leads to a decline in mutual funds' local overweighting by lowering information-processing costs. If geographic proximity provides an advantage partly because local firms are easier to evaluate, then improvements in information-processing technology should disproportionately reduce the relative cost of evaluating distant firms and narrow the local–distant portfolio-weight gap. We first estimate the baseline effect of ChatGPT's release on the local–distant portfolio-weight gap and examine its dynamics over time. We then decompose the adjustment into extensive and intensive margins to assess how each contributes to the decline in local bias. Finally, we explore heterogeneity across firm and fund characteristics to provide further evidence on the information-processing mechanism.

3.1 Baseline results

We examine whether the emergence of GenAI reshapes the geographic composition of mutual fund portfolios. Specifically, we ask whether funds rebalance away from locally headquartered firms toward more distant ones once GenAI becomes widely available.

Our empirical strategy follows a difference-in-differences design that compares funds' portfolio weights in local and distant stocks before and after the ChatGPT launch in November 2022. The analysis is conducted at the fund-firm-quarter level. Formally, we estimate the following specification:

$$Weight_{i,j,t} = \alpha_i + \lambda_j + \delta_t + \beta_1 Local_{i,j} + \beta_2 \left(Local_{i,j} \times PostGPT_t\right) + \boldsymbol{\gamma}' \mathbf{x}_{i,t-1} + \boldsymbol{\theta}' \mathbf{z}_{j,t-1} + \varepsilon_{i,j,t}, \quad (2)$$

where $Weight_{i,j,t}$ is the portfolio weight of stock j in fund i's portfolio at quarter t. $Local_{i,j}$ is an indicator equal to one if firm j's headquarters is located within 100 miles of fund i's headquarters, and zero otherwise. In these baseline tests, we focus on the six calendar quarters before and after the introduction of ChatGPT in November 2022. $PostGPT_t$ equals one for the quarters after the ChatGPT launch and zero for the quarters before the shock. Since mutual fund holdings are ob-

served at a quarterly frequency, the ChatGPT-launch quarter ($2022Q4$) is only partially exposed to the shock and cannot be cleanly classified as either a pre- or post-GPT period. We therefore exclude this transition quarter from the sample, ensuring a clean comparison between pre- and post-event portfolio allocations. The vectors $\mathbf{x}_{i,t-1}$ and $\mathbf{z}_{j,t-1}$ contain lagged fund-level and firm-level controls measured at the end of the quarter $t-1$, respectively.

The specification includes fund fixed effects, α_i, stock fixed effects, λ_j, and year-quarter fixed effects, δ_t. Fund fixed effects absorb time-invariant differences across mutual funds, such as investment style and portfolio strategy. Stock fixed effects absorb time-invariant differences across stocks that may affect their average portfolio weights across funds. Year-quarter fixed effects control for aggregate shocks common to all funds and stocks in a given quarter, such as market-wide conditions and macroeconomic shocks. As a result, identification comes from within-fund variation in portfolio weights on local relative to non-local stocks around the introduction of ChatGPT.

The coefficient β_1 captures the average difference in portfolio weights between local and distant stocks during the pre-GenAI period. The coefficient of interest, β_2, measures how this gap changes after the introduction of ChatGPT[20]. A negative β_2 indicates that the relative overweighting of local stocks declines following the ChatGPT launch, consistent with the interpretation that GenAI lowers the information-processing barriers associated with geographic distance and thereby reduces local investment bias. Standard errors are clustered at the fund level.

We report the results of estimating Equation 2 in Table 3. Column 1 shows that the coefficient on $Local_{i,j}$ is positive and statistically significant, indicating that funds, on average, assign higher portfolio weights to geographically proximate stocks during the pre-GenAI period. This pattern is consistent with prior studies documenting that mutual fund managers exhibit a systematic preference for nearby firms (see Coval and Moskowitz, 1999; Coval and Moskowitz, 2001; Baik, Kang and Kim, 2010; Bernile, Kumar and Sulaeman, 2015). The coefficient on the interaction term $Local_{i,j} \times PostGPT_t$ provides the main evidence for the change in local bias after the ChatGPT launch.

[20]As GenAI adoption diffused gradually and heterogeneously across asset managers, our $PostGPT_t$ indicator should be interpreted as capturing the onset of a new information environment rather than instantaneous universal adoption. Our difference-in-differences estimates therefore correspond to "intention-to-treat" effects, with further heterogeneity analyses in Section 3.4 exploiting predicted variation in adoption and usefulness of GenAI.

Its estimate indicates a statistically significant decline in the local-versus-distant portfolio weight gap, suggesting that funds reduce their relative overweighting of nearby stocks after the ChatGPT launch.

Column 2 in Table 3 examines the robustness of this finding by including fund- and firm-level covariates. The fund-level controls include fund size, returns, net flows, turnover ratio, expense ratio, and an indicator for whether the fund primarily targets institutional investors. The firm-level controls include firm size, book-to-market ratio, profitability, leverage, and price momentum. The estimated coefficient on the interaction term remains negative and statistically significant, indicating that the decline in local overweighting is not driven by observable fund- or firm-level characteristics.

We further verify that the results continue to hold under two alternative sets of fixed-effect structures. In Columns 3 and 4 of Table 3, we control fund × industry, stock and year-quarter fixed effects. Fund × industry fixed effects absorb time-invariant differences in each fund's exposure to particular industries, such as those driven by their investment mandates, style preferences, or industry specialization. This also helps ensure that the estimated decline in local overweighting is identified from variation in portfolio weights between local and distant stocks within the same fund-industry category.

In Columns 5 and 6 of Table 3, we control fund × firm and year-quarter fixed effects, as in Kuchler et al. (2022). These fixed effects absorb all time-invariant heterogeneity at the fund-firm pair level, including a fund's persistent tendency to hold a particular stock, long-run familiarity with the firm, and any stable geographic or informational relationship between the fund and the firm. Therefore, the the coefficients of interest are identified from whether the change in a given fund's weight in a given stock around the ChatGPT launch differs systematically depending on whether the fund-firm pair is local or distant.

Across both sets of fixed effect structures, the estimated coefficient on $Local_{i,j} \times PostGPT_t$ remains negative and statistically significant. More generally, the stability of the estimates of β_2 as we progressively add controls and more demanding fixed effects alleviates concerns that the results are driven by omitted unobserved confounders.

Are the results mechanically driven by funds' industry reallocation or over-performance of distant firms?

A potential concern is that the estimated decline in local overweighting may mechanically reflect a post-ChatGPT increase in the portfolio weights of technology stocks, arising either from active industry reallocation or from relative stock-price appreciation. This concern is particularly relevant to our setting because many mutual funds are clustered in New York City and Boston, which together account for 26% of the funds in our sample, whereas technology firms are disproportionately concentrated in West Coast hubs such as the San Francisco Bay Area and Seattle. This geographic pattern creates two related concerns. First, if funds headquartered in New York City and Boston increase their exposure to technology firms after the ChatGPT launch because these firms potentially benefit from AI-related attention and growth opportunities, then their weights of geographically distant stocks would rise mechanically. Second, even absent active reallocation, the potential post-ChatGPT appreciation of technology stock prices could mechanically raise the portfolio weights of distant holdings. In both cases, the estimated decline in local overweighting could be driven by a post-GPT rotation toward, or value appreciation, in technology stocks rather than a broader reduction in funds' reliance on geographic proximity.

We address these possibilities in two ways. First, we re-estimate the baseline specification in Equation 2 after excluding observations that are most likely to be relevant to this concern. Specifically, we exclude funds headquartered in New York and Boston and technology firms from the sample, respectively. We report the results in Panel A of Table 4. The estimated coefficient on $Local_{i,j} \times PostGPT_t$ remains negative and statistically significant, indicating that the decline in local overweighting is not driven by a post-GenAI shift toward technology firms. Second, we augment the baseline specification by controlling for contemporaneous stock returns, thereby accounting for the possibility that changes in portfolio weights mechanically reflect valuation effects rather than active portfolio reallocation. The results are reported in Panel B of Table 4. We find our results persist, suggesting that the decline in local overweighting is not merely a mechanical consequence

of distant stock price appreciation[21].

More generally, however, this concern is not limited to technology firms or to funds located in New York and Boston. A fund's industry allocation may evolve over time, and such changes can affect measured local overweighting if industries are geographically clustered. To address this broader concern, we employ more restrictive fixed-effect structures. In Column 1 of Table 5, in addition to fund × industry and stock fixed effects, we further introduce fund × year-quarter and industry × year-quarter fixed effects. These high-dimensional fixed effects absorb time-varying heterogeneity at both the fund and industry levels. In Column 2, we replace industry × year-quarter fixed effects with stock × year-quarter fixed effects, thereby controlling stock-specific shocks in each quarter. In column 3, we include fund × industry × year-quarter and stock fixed effects. These fixed effects absorb each fund's exposure to a given industry in each quarter, allowing a fund's industry allocation to vary flexibly quarter by quarter.

Across these specifications, the estimated coefficient on $Local_{i,j} \times PostGPT_t$ remains negative and statistically significant. The estimates are also quantitatively similar to the baseline results, confirming that the decline in the relative overweighting of local stocks after the ChatGPT launch is robust to these saturated fixed-effect structures. These results further assure that our finding is not explained by quarter-specific fund behavior, industry- or stock-level shocks, or changes in funds' industry allocations over time.

3.2 Dynamics of local overweighting around the ChatGPT launch

We next examine the dynamics of local overweighting around the ChatGPT launch. This analysis serves two purposes. First, it allows us to assess whether portfolio weights on local and distant stocks follow parallel trends prior to the introduction of ChatGPT—a key identifying assumption of our difference-in-differences design. Second, it reveals whether the reduction in local overweighting

[21]We interpret these results with caution. As stock returns may themselves be affected by the introduction of ChatGPT, controlling for contemporaneous returns in the baseline specification risks a "bad control" problem that could confound the estimated treatment effect (see Caetano et al., 2024). For this reason, we do not include contemporaneous returns in the baseline regression, but instead report this specification as an additional robustness check.

emerges immediately following the shock or unfolds gradually over subsequent quarters. Formally, we estimate the following specification:

$$Weight_{i,j,t} = \alpha_i + \lambda_j + \delta_t + \beta_1 Local_{i,j} + \sum_{\tau \neq -1} \beta_\tau \left(\mathbb{1}_{t-T_0=\tau} \times Local_{i,j} \right) + \varepsilon_{i,j,t}, \tag{3}$$

where $\mathbb{1}_{t-T_0=\tau}$ is an event-time indicator equal to one when quarter t is τ quarters away from the ChatGPT-launch quarter, $T_0 = 2022Q4$. All other variables are defined as before. The dynamic coefficients, β_τ, measure how the local-versus-distant portfolio weight gap in event quarter τ differs relative to the omitted reference quarter, $\tau = -1$, the quarter immediately preceding the GenAI shock.

Table 6 presents the results of regression 3. The four columns vary in the inclusion of covariates and fixed-effect structures, but yield broadly consistent estimates. Focusing first on the pre-event period, the coefficients from $\tau = -6$ to $\tau = -2$ are statistically indistinguishable from zero across all specifications. This suggests that portfolio weights for local and distant stocks do not exhibit differential trends before the ChatGPT launch, lending support to the parallel-trends assumption. Turning to the post-event period, the estimates reveal a gradual decline in local overweighting. The coefficient at $\tau = +1$ is close to zero and statistically insignificant, indicating that funds do not immediately adjust their relative allocation to local stocks in the quarter following the ChatGPT launch. Starting from $\tau = +2$, the coefficients turn negative and begin to gain statistical significance, and from $\tau = +3$ onward they are consistently negative and statistically significant across all specifications, with magnitudes that grow steadily over time.

Figure 3 visualises the dynamic pattern of local overweighting based on the estimates from column 1. The pattern is consistent with the gradual diffusion of GenAI into mutual fund investment workflows. While the initial launch of ChatGPT in November 2022 represents the first broad GenAI availability shock, fund managers likely needed time to experiment with these tools and incorporate them into information acquisition and portfolio decision-making processes. The estimated effect becomes stronger after the release of GPT-4 in March 2023, which marked a substantial improvement of model capability in dealing with complex text-based tasks. The effect deepens further around the release of GPT-4 Turbo in November 2023, which expanded the context window and reduced usage costs, making large-scale processing of firm-level information more feasible. This timing suggests

that successive improvements in the accessibility, capability, and scalability of GenAI progressively eroded the information-processing barriers associated with evaluating geographically distant firms.

Overall, the dynamic evidence reinforces the baseline results. The absence of pre-trends between local and distant portfolio weights supports the validity of the empirical design, while the gradual post-shock decline in local overweighting suggests that the effect strengthens as GenAI becomes more capable and is gradually incorporated into fund managers' investment processes.

3.3 Extensive and intensive margins of portfolio reallocation

GenAI may reshape funds' holdings of local versus distant stocks in different ways. It may expand the set of firms that funds are willing to consider, leading them to enter more distant stocks or exit from local stocks. Alternatively, it may affect how funds allocate capital across existing holdings, leading them to reduce local positions and increase distant positions without necessarily changing the set of stocks they hold. To understand the channel behind the decline in local investment bias, we examine whether the effect operates through changing the decisions to hold particular stocks (extensive margin), through the rebalancing of positions that funds already hold (intensive margin) or through both channels.

More specifically, we first examine whether the probability that funds hold local versus distant stocks changes after the launch of ChatGPT. We define a holding indicator, $Held_{i,j,t}$, equal to one if fund i holds stock j in quarter t, and zero otherwise, and then re-estimate equation 2 with $Held_{i,j,t}$ as the dependent variable. The results are reported in Column 1 of Table 7. The coefficient on $Local_{i,j} \times PostGPT_t$ is negative and statistically significant, suggesting that funds reduce their propensity to hold local stocks relative to distant ones following the launch of ChatGPT. This points to a weakening of local investment bias along the extensive margin.

A natural follow-up question is whether the decline in the relative probability of holding local stocks is driven by greater entry into distant stocks relative to local ones, greater exit from local holdings relative to distant ones, or both. We therefore examine the transition dynamics underlying changes in the holding decision. In the spirit of He, Hu and Li (2026), we define $ExtensiveIn_{i,j,t}$ as

an indicator equal to one when a fund initiates a position in a stock that it did not hold in the previous quarter; the corresponding regression is estimated on the sample of fund-stock pairs that were not held in the previous quarter and are therefore at risk of entry. Analogously, $ExtensiveOut_{i,j,t}$ is an indicator equal to one when a fund fully liquidates a position that it held in the previous quarter, and the corresponding regression is estimated on fund-stock pairs that were held in the previous quarter and are thus at risk of exit.

Column 2 and 3 of Table 7 report the results from estimating Equation 2 after replacing the dependent variable with $ExtensiveIn_{i,j,t}$ and $ExtensiveOut_{i,j,t}$, respectively. The coefficient on $Local_{i,j} \times PostGPT_t$ is negative and statistically significant in the entry specification, but close to zero and statistically insignificant in the exit specification. This suggests that the decline in local holding probability is driven primarily by funds initiating more new positions in distant stocks relative to local stocks, rather than by funds disproportionately unwinding existing local holdings. The insignificant exit result is also informative. It suggests that GenAI does not fully eliminate the value of local information advantages. For local stocks that funds already hold, managers may have accumulated firm-specific knowledge and built local information networks through prior research. As a result, even if GenAI lowers the cost of evaluating distant firms, existing local information advantages may continue to make local holdings valuable.

Finally, we examine whether the decline in local bias also operates through the intensive margin, which captures the adjustments to a fund's continuing holdings through additional buying or selling. Specifically, we restrict the sample to fund-stock pairs that are held in both the previous and current quarters and use the change in portfolio weight $\Delta Weight_{i,j,t}$ as the dependent variable. The results are reported in column 4 of Table 7. The coefficient on $Local_{i,j} \times PostGPT_t$ is negative and statistically significant, indicating that, , after the launch of ChatGPT, funds reallocate portfolio weight from existing local holdings toward existing distant holdings.

Overall, our results suggest that the decline in local bias operates through two channels. First, funds initiate more positions in distant stocks relative to local stocks. Second, within continuing holdings, they rebalance away from local stocks and toward distant ones. At the same time, funds do

not simply abandon existing local holdings outright, consistent with the view that local information advantages may still retain some value. Rather, GenAI appears to erode the information frictions traditionally associated with geographic remoteness. By making distant firms easier to evaluate, it reduces funds' reliance on geographic proximity when deciding which new stocks to buy and how to rebalance positions they already hold.

3.4 Cross-sectional variation in benefits from GenAI

Our results so far establish that mutual funds reduce their local overweighting following the launch of ChatGPT. We argue that this pattern reflects a decline in the information-processing frictions associated with evaluating geographically distant firms. We next examine heterogeneity in the post-ChatGPT decline in local overweighting along firm and fund characteristics to provide further evidence on this mechanism.

3.4.1 Firm-level disclosure complexity

If the reduction in local bias is driven by GenAI lowering the cost of processing firm-level disclosures, the effect should be more pronounced for firms whose disclosures are difficult to read and interpret. GenAI is particularly useful for parsing long and linguistically complex text, so these firms—which impose the greatest information-processing burden on investors—stand to benefit most from GenAI-assisted analysis. By contrast, for firms with short and easier-to-interpret filings, the marginal value of GenAI is limited, and the post-ChatGPT change in local bias should accordingly be muted.

We test this prediction using two measures constructed from firms' 10-K filings, each capturing a different dimension of information-processing costs. The first, filing size, proxies for the volume of information investors must process: longer filings require more time and effort to read, search, and summarize. The second, the SMOG index, summarizes overall readability by combining sentence structure with word difficulty, with higher values indicating that greater reading effort is required to understand the underlying information. We obtain both measures from the SEC An-

alytics Suite provided by WRDS. For each measure, we compute preshock processing costs as the average value over the pre-ChatGPT period. We then split firms into low- and high-processing-cost subsamples based on the median of this preshock average and re-estimate the baseline specification separately within each group.

Table 8 Panel A reports the results. From Columns 1 and 3, we observe that, for firms in the low-processing-cost subsamples (i.e. smaller filing size, or a lower SMOG score), the estimated coefficient on $Local_{i,j} \times PostGPT_t$ is small and statistically indistinguishable from zero. By comparison, for firms in the high-processing-cost subsamples, the estimated coefficient on the interaction term is consistently negative and statistically significant at the 1% level, as shown in Columns 2 and 4. These results indicate that the post-ChatGPT decline in local overweighting is concentrated among firms with more complex textual disclosures.

Since firm size is positively correlated with disclosure complexity, we further examine whether disclosure complexity explains the post-ChatGPT decline in local overweighting beyond its correlation with firm size. We perform an independent two-way sort that classifies firms by size and disclosure complexity, and re-estimate the baseline specification across the resulting size–complexity groups. The results are reported in Panel B of Table 8. For both complexity measures, the decline in local overweighting is concentrated among large firms with high disclosure complexity, while the estimated effect is close to zero among large firms with low disclosure complexity. Notably, we do not find evidence of a decline in local overweighting among small firms, regardless of disclosure complexity—likely because pre-ChatGPT local overweighting is itself concentrated among large firms, leaving little scope for the bias to decline further among small firms in the first place. Overall, conditional on firm size, the decline in local overweighting remains concentrated among firms with more complex disclosures, suggesting that our complexity results are not solely attributable to the correlation between disclosure complexity and firm size.

Taken together, these findings support the interpretation that the emergence of GenAI weakens local investment bias by reducing information-processing costs, particularly when firms' disclosures are lengthy or hard to interpret and therefore therefore impose greater demands on investors'

processing capacity.

3.4.2 Quantitative versus discretionary funds

We next examine whether the decline in local overweighting varies with the extent to which funds rely on managerial discretion in their investment decisions. If GenAI complements human intelligence by effectively augmenting managers' capacity to process information, its effect should be stronger when human capital plays an important role in making investment decisions. Building on Abis (2022), we employ ChatGPT 5.5 to analyse the Principal Investment Strategy (PIS) section of fund prospectuses[22] and classify funds as quantitative if statistical, algorithmic or other model-based methods determine security selection and portfolio construction, and as discretionary if human judgment is indispensable to investment decisions.

Column 1 of Table 9 reports the results for funds that rely purely on quantitative investment strategies. The coefficient on $Local_{i,j}$ is statistically indistinguishable from zero for quantitative funds, suggesting that these funds do not exhibit a clear local overweighting pattern before the ChatGPT launch. The coefficient on $Local_{i,j} \times PostGPT_t$ is negative but statistically insignificant, indicating that there is no statistically detectable change in local overweighting for purely quantitative funds after the GenAI shock. The results for discretionary funds are shown in Column 2. The coefficient on $Local_{i,j}$ is positive and statistically significant, indicating that discretionary funds overweight geographically proximate stocks before the ChatGPT launch. More importantly, the coefficient on $Local_{i,j} \times PostGPT_t$ is negative and statistically significant, suggesting that this local overweighting declines after the GenAI shock.

This contrast is economically intuitive. Quantitative funds already rely heavily on systematic signals and model-based decision rules, which may leave less room for local familiarity or soft information to shape portfolio allocations. As a result, their baseline local bias is weaker, and there is less scope for GenAI to further reduce it. By contrast, discretionary funds rely more

[22]Data on funds' investment strategies come from the SEC's "Mutual Fund Prospectus Risk/Return Summary Data Sets", which are publicly available on the SEC's website. See: https://www.sec.gov/data-research/sec-markets-data/mutual-fund-prospectus-riskreturn-summary-data-sets.

on human judgment for information processing. These funds are more likely to face attention constraints and to benefit from tools that reduce the cost of evaluating distant firms. The results are therefore consistent with the interpretation that GenAI weakens local bias by reducing the information-processing constraints faced by human decision-makers.

3.4.3 Fund managerial structure

Kacperczyk, Van Nieuwerburgh and Veldkamp (2016) argue that fund managers face limited information-processing capacity, but this capacity can be expanded through analytical resources such as additional analyst hires. GenAI plausibly serves an analogous function, expanding a manager's effective capacity to process firm-level information. Its value should therefore be greater for managers who are more capacity-constrained *ex ante*.

Fund managerial structure provides a natural setting to test this implication. Funds managed by larger teams may already have greater internal processing capacity through specialisation, division of labour, and information sharing. By contrast, funds managed by fewer named managers are likely to face tighter attention constraints. We therefore split the sample into team-managed and non-team-managed funds, following CRSP's classification[23], to test whether the effect of the emergence of GenAI on reducing local overweighting is stronger when managerial capacity constraints are more binding.

Table reports the results. Column 1 presents the estimates for team-managed funds, while Column 2 reports the results for non-team-managed funds. The coefficient on $Local_{i,j}$ is positive in both subsamples, but it is statistically significant only for non-team-managed funds. This suggests that, before the introduction of ChatGPT, local overweighting is more pronounced among non-team-managed funds and is not statistically evident among team-managed funds. This is consistent with the idea in Van Nieuwerburgh and Veldkamp (2010) that funds with more limited information capacity rely more heavily on geographically proximate firms, for which information is easier to learn and interpret.

[23]A fund is classified as team-managed when the CRSP manager-name field (*mgr_name*) is reported as "team managed", and as non-team-managed otherwise.

More importantly, the coefficient on $Local_{i,j} \times PostGPT_t$ is negative in both subsamples, but it is larger in magnitude and statistically significant only among non-team-managed funds. For team-managed funds, the coefficient is -0.008 and statistically insignificant. By contrast, for non-team-managed funds, the coefficient is -0.019 and statistically significant at the 1% level. Taken together, these results are consistent with the view that GenAI relaxes information-processing constraints and is especially valuable when managerial information-processing capacity is more limited.

4 The informativeness of portfolio reallocation

We argue that GenAI facilitates the processing of firm disclosures and other public information, thereby attenuating the information asymmetry that has traditionally disadvantaged distant investors. If information processing constitutes the operative channel, this should manifest not merely in the reallocation of portfolio weights toward distant firms, but also in the informativeness of that reallocation. That is, managers possessing more precise firm-specific signals about distant firms should assign larger portfolio weights to those that subsequently experience positive fundamental shocks and smaller weights to those that experience negative shocks. In this section, we therefore examine whether the post-ChatGPT portfolio reallocation is associated with improved investment outcomes for distant holdings.

4.1 Fund managers' stock-picking ability

First, we examine fund managers' stock-picking ability as a direct test of whether the post-ChatGPT decline in local overweighting is accompanied by improved selection among distant stocks, as opposed to uninformed portfolio reallocation toward previously underweighted stocks. We focus on stock-picking ability because the hypothesised channel operates primarily through the processing of firm-specific fundamental information.

For each fund, we calculate the stock-picking ability separately for its local and distant holdings. Following Kacperczyk, Van Nieuwerburgh and Veldkamp (2016), the stock-picking ability for fund

i, portfolio segment $g \in \{Local, Distant\}$ and quarter t, is defined as:

$$StockPicking_{i,g,t} = \frac{1}{N_{i,g,t}} \sum_{j \in \mathcal{G}_{i,t}^{g}} \left(w_{i,j,t} - w_{m,j,t}\right) \left(R_{j,t+1} - \beta_{j,t} R_{m,t+1}\right), \tag{4}$$

where $\mathcal{G}_{i,t}^{g}$ denotes the set of stocks in fund i's portfolio segment g at quarter t, and $N_{i,g,t}$ is the number of stocks in that segment. $w_{i,j,t}$ is fund i's portfolio weight in stock j at quarter t, and $w_{m,j,t}$ is stock j's market weight at the same time. Thus, $w_{i,j,t} - w_{m,j,t}$ captures the fund's active weight in stock j relative to the market portfolio. $R_{j,t+1}$ is the realised return of stock j over the subsequent quarter. The term $\beta_{j,t} R_{m,t+1}$ captures the systematic component of stock j's subsequent return, where $\beta_{j,t}$ is estimated using a 12-month rolling-window regression of stock j's excess returns on market excess returns from month $t-11$ to month t. The difference $R_{j,t+1} - \beta_{j,t} R_{m,t+1}$ thereby captures the idiosyncratic component of the subsequent return, which is unknown at the time of portfolio formation. Intuitively, $StockPicking_{i,g,t}$ is higher when a fund actively overweights stocks that subsequently realise high idiosyncratic returns and underweights stocks that subsequently realise low idiosyncratic returns. Therefore, the measure captures the informativeness of a fund's security selection within its local and distant portfolios.

To test whether the launch of ChatGPT affects fund managers' stock-picking ability across local and distant holdings, we estimate the following linear regression model:

$$StockPicking_{i,g,t} = \alpha_i + \delta_t + \theta_1 Local_g + \theta_2 \left(Local_g \times PostGPT_t\right) + \boldsymbol{\theta}' \mathbf{x}_{i,t} + \varepsilon_{i,g,t}, \tag{5}$$

where i indexes funds, $g \in \{Local, Distant\}$ indexes the portfolio segments, and t indexes year-quarters. $StockPicking_{i,g,t}$ is the stock-picking measure defined above. $Local_g$ is an indicator equal to one for the local portfolio segment and zero for the distant segment. $PostGPT_t$ equals one for quarters after the release of ChatGPT. The coefficient of interest is θ_2, which captures the change in stock-picking ability of local portfolios relative to distant portfolios after the ChatGPT launch. $\mathbf{x}_{i,t}$ is a vector of fund characteristics, including fund flows, turnover ratio, expense ratio, an institutional-fund indicator, and fund size. α_i and δ_t denote fund and year-quarter fixed effects, respectively. Alternatively, we replace fund and year-quarter fixed effects with style $\times$ year-quarter fixed effects to absorb time-varying shocks common to funds with similar investment

styles. Standard errors are clustered at the fund level.

The results are reported in Table 11. Across all specifications, the coefficient on $Local_g$ is positive and statistically significant, indicating that before the ChatGPT launch, local portfolios exhibit higher stock-picking ability than distant portfolios. More importantly, the coefficient on $Local_g \times PostGPT_t$ is negative and statistically significant across columns. This suggests that, after the ChatGPT launch, the stock-picking advantage of local portfolios relative to distant portfolios declines. The result is consistent with GenAI reducing the information-processing disadvantage associated with distant firms, making allocations to them more informative.

4.2 Local versus distant portfolio performance

We next examine whether the relative improvement in the informativeness of distant allocations translates into portfolio performance. If GenAI reduces the information-processing disadvantage associated with distant firms, the performance of distant portfolios should improve relative to that of local portfolios. We therefore construct local and distant portfolio returns for each fund-quarter and estimate whether the performance gap between local and distant holdings changes after the introduction of ChatGPT.

Following Pool, Stoffman and Yonker (2012), we define the return of portfolio segment $g \in \{Local, Distant\}$ for fund i, based on its holdings at the end of quarter t, as:

$$R_{i,g,t} = \sum_{j \in \mathcal{G}_{i,t}^{g}} \left(\frac{w_{i,j,t}}{\sum_{j \in \mathcal{G}_{i,t}^{g}} w_{i,j,t}} \right) R_{j,t+1}, \tag{6}$$

where $\mathcal{G}_{i,t}^{g}$ denotes the set of stocks in fund i's portfolio segment g at quarter t, $w_{i,j,t}$ is the actual portfolio weight of fund i in stock j at the end of quarter t, and R_j is the realised excess return of stock j. We use two different excess return measures for R_j. The first is the stock's excess return over the risk-free rate. The second is the DGTW-adjusted return (Daniel et al., 1997), which is calculated by subtracting from each stock's return the return of a benchmark portfolio formed on size, book-to-market, and momentum characteristics[24].

[24]The DGTW benchmark-adjustment procedure sorts stocks into characteristic-based portfolios using size, book-to-market, and momentum. Specifically, stocks are first sorted into size quintiles; within each size quintile, they are then

To test whether ChatGPT launch changes the relative performance of local and distant holdings, we estimate the follwing regression model:

$$R_{i,g,t} = \alpha_i + \delta_t + \varphi_1 Local_g + \varphi_2 \left(Local_g \times PostGPT_t\right) + \boldsymbol{\lambda}' \mathbf{x}_{i,t} + \varepsilon_{i,g,t}, \tag{7}$$

where $R_{i,g,t}$ denotes the local or distant portfolio return defined above. All other variables are defined as in Equation 5. The coefficient of interest is φ_2, which captures the change in the performance of local holdings relative to distant holdings after the introduction of ChatGPT. Standard errors are clustered at the fund level.

Table 12 reports the regression results. The coefficient on $Local_g \times PostGPT_t$ is negative and statistically significant across all specifications and for both return measures, indicating that distant portfolio performance improves relative to local portfolio performance in the post-ChatGPT period. Specifically, the point estimates in Columns 1 to 3 indicate that excess returns on local portfolios are approximately $1.56 - 1.60$ percentage points lower than those on distant portfolios after ChatGPT launch, relative to the pre-ChatGPT period. Columns 4 to 6 show a similar pattern using DGTW-adjusted returns, with local portfolios underperforming distant portfolios by approximately $0.31 - 0.33$ percentage points after ChatGPT launch, relative to the pre-ChatGPT period.

Overall, these results complement the stock-picking evidence by showing that the relative improvement in the informativeness of distant allocations is also reflected in realised portfolio performance. The findings are consistent with the interpretation that ChatGPT helps fund managers process firm-specific information about distant stocks, thereby enhancing the investment decisions in them.

5 Robustness

We conduct a series of robustness tests.

sorted into book-to-market quintiles; and within each size–book-to-market quintile, they are further sorted into momentum quintiles. This triple-sort procedure produces 125 characteristic-based benchmark portfolios. The benchmark return for each portfolio is calculated as the value-weighted average return of the stocks in that portfolio. The DGTW-adjusted return is then obtained by subtracting from each stock's return the return of its matched benchmark portfolio.

5.1 Alternative definition of local firms

In the baseline analysis, we define a fund-firm pair as local if the distance between the fund's headquarters and the firm's headquarters is less than 100 miles. To ensure that our results are not driven by this particular cutoff, we repeat the baseline specification using alternative distance thresholds to define local fund-firm pairs. Specifically, we define $Local_{ij}$ using cutoffs of 150, 200, and 250 miles. As shown in Table 13, the estimated coefficient on $Local_{i,j} \times PostGPT_t$ remains negative and statistically significant across these alternative definitions. These results indicate that the documented decline in local overweighting after the ChatGPT launch is not sensitive to the particular distance threshold used to classify local holdings.

5.2 Alternative holdings measure

We examine whether the results are sensitive to the way that fund holdings are measured. In the baseline analysis, we use portfolio weights as the dependent variable, which captures the economic importance of each stock in a fund's portfolio. As an alternative measure, we re-estimate the baseline specification using the logarithm of the number of shares of a given stock held by the fund. This measure captures changes in the quantity of shares held and is less affected by changes in stock price or portfolio values. From Table 14, we observe that the estimated coefficient on $Local_{i,j} \times PostGPT_t$ remains negative and statistically significant.

5.3 Alternative event window

In our baseline specification, we use an event window of $[-6Q, +6Q]$ around ChatGPT's release. To ensure that our results are not driven by this particular choice, we re-estimate the baseline specification using alternative event windows of $[-8Q, +8Q]$ and $[-4Q, +4Q]$. The results are reported in Table 15. The wider$[-8Q, +8Q]$ window yields a slightly larger effect than the baseline specification, whereas the narrower $[-4Q, +4Q]$ window produces a somewhat smaller effect. This pattern is consistent with the event-study evidence that the decline in local overweighting

develops gradually following ChatGPT's release and becomes more pronounced in later quarters. Overall, the coefficient on $Local_{i,j} \times PostGPT_t$ remains negative and statistically significant in most specifications, indicating that our finding is not sensitive to the choice of event window.

5.4 Alternative clustering of standard errors

We examine whether our inference is sensitive to alternative clustering of standard errors. In the baseline specification, standard errors are clustered at the fund level. We repeat the analysis by clustering standard errors at the stock level and by using two-way clustering at the fund and stock levels. The estimated coefficient on $Local_{i,j} \times PostGPT_t$ in Table 16 remains negative and statistically significant under both approaches.

5.5 Controlling for pre-event covariates

As an additional robustness test, we replace the time-varying fund- and stock-level covariates in the baseline specification with their pre-event values interacted with the post-event indicator, *PostGPT*. This addresses the concern that these covariates may themselves be affected by the emergence of GenAI and therefore give rise to a potential "bad control" problem (see Caetano et al., 2024). As shown in Table 17, the coefficient on $Local_{i,j} \times PostGPT_t$ remains negative and statistically significant across all specifications, with estimates ranging from -0.012 to -0.016. The similarity in magnitude to our baseline estimates indicates that our findings are not sensitive to the inclusion of potentially endogenous post-event covariates.

5.6 Placebo definition of local firms

We conduct placebo tests using alternative definitions of local fund-firm pairs. Specifically, instead of classifying geographically proximate pairs as local, we define placebo local pairs using random distant distance bands: $900 - 1,000$ miles, and $1,400 - 1,500$ miles. These distance bands are unlikely to capture meaningful local informational advantages. Therefore, if the baseline

results reflect a genuine decline in local overweighting driven a reduction in the relative informational advantage of local firms over distant firms, we should not observe a significant effect under these placebo definitions. Consistent with this expectation, the estimated coefficient on $PlaceboLocal_{i,j} \times PostGPT_t$ in Table 19 is statistically indistinguishable from zero across these placebo definitions for local holdings.

5.7 Placebo event date

As another set of placebo tests, we assign three pseudo-event dates—July 1, 2018; January 1, 2019; and July 1, 2019—that predate the emergence of GenAI and fall within a period without comparable advances in information-processing technology[25], and re-estimate the baseline specification. Across all placebo tests in Table 19, the coefficient on $Local_{i,j}$ remains positive and statistically significant, indicating that mutual funds consistently overweight geographically proximate stocks during these earlier periods. In contrast, the estimated coefficients on $Local_{i,j} \times PlaceboPost_t$ are economically small and statistically insignificant. The absence of significant effects around these placebo dates suggests that our baseline findings are unlikely to be explained by pre-existing differential trends between local and distant holdings or slow-moving systematic changes in portfolio allocations.

6 Conclusion

The underlying causes of local investment bias have long been of interest to finance researchers. Existing explanations emphasise either informational advantages associated with geographic proximity or behavioural preferences for nearby firms. Within the informational-advantage view, however, considerably less is known about whether local bias arises because local investors face lower costs of processing public information about nearby firms. This paper examines this under-explored margin by exploiting the emergence of GenAI as a shock that reduces information-processing frictions.

Using a difference-in-differences strategy, we show that mutual funds significantly reduce their

[25]We restrict the placebo dates to the pre-COVID period because the pandemic generated substantial disruptions to mobility and face-to-face information acquisition, which could independently affect local investment bias (see Lee, 2026).

local overweighting following the emergence of GenAI. Before the release of ChatGPT, funds systematically allocate greater portfolio weights to geographically proximate firms. Following its release, this local–distant allocation gap declines substantially. The adjustment occurs through both the extensive and intensive margins: funds become more willing to initiate positions in distant firms and, among continuing holdings, reallocate portfolio weight away from local firms toward distant firms. The effect is concentrated in settings where information-processing demands are greater. In particular, the decline in local overweighting is more pronounced among holdings of firms with more complex textual disclosures, consistent with GenAI lowering the costs of extracting, interpreting, and synthesising firm-level information. The effect is also evident among funds where information-processing constraints are more likely to bind, with statistically significant declines among discretionary and non-team-managed funds.

The portfolio reallocation is also informative. Prior to the emergence of GenAI, local holdings exhibit stronger stock-picking performance than distant holdings, consistent with an informational advantage associated with geographic proximity. Following ChatGPT's release, this local stock-picking advantage declines, as the informativeness of distant holdings improves relative to that of local holdings. Consistent with this evidence, the performance gap between local and distant portfolio segments also narrows. These findings suggest that GenAI does not simply alter where funds invest, but may also improve the quality of investment decisions involving geographically distant firms.

Our findings have broader implications for understanding local investment bias. Geographic proximity may matter not only because nearby investors have access to different information, but also because they can process and interpret available information at lower cost. By reducing these processing costs, GenAI can weaken the informational advantage associated with proximity and broaden the set of firms that investors can evaluate effectively. More generally, our results shed light on why geographic proximity continues to shape investment decisions in an increasingly information-rich environment and show how advances in information-processing technology can reshape persistent portfolio distortion stemming from investors' limited attention and analytical capacity.

References

Abis, S. 2022. "Man vs. machine: Quantitative and discretionary equity management." Unpublished Paper.

Alok, S, N Kumar, and R Wermers. 2020. "Do Fund Managers Misestimate Climatic Disaster Risk." *The Review of Financial Studies*, 33(3): 1146–1183.

Bai, J, and M Massa. 2026. "Is human-interaction-based information substitutable?" *Journal of Financial Intermediation*, 67: 101210.

Baik, B, J.-K Kang, and J.-M Kim. 2010. "Local institutional investors, information asymmetries, and equity returns." *Journal of Financial Economics*, 97(1): 81–106.

Barber, B, and T Odean. 2002. "Online investors: Do the slow die first?" *Review of Financial Studies*, 15(2): 455–487.

Ben-David, I, J Birru, and A Rossi. 2019. "Industry familiarity and trading: Evidence from the personal portfolios of industry insiders." *Journal of Financial Economics*, 132(1): 49–75.

Bernile, G, A Kumar, and J Sulaeman. 2015. "Home away from home: Geography of information and local investors." *Review of Financial Studies*, 28(7): 2009–2049.

Bernstein, S, X Grioud, and R. R Townsend. 2016. "The impact of venture capital monitoring." *Journal of Finance*, 71(4): 1591–1622.

Bertomeu, J, Y Lin, Y Liu, and Z Ni. 2025. "The impact of generative AI on information processing: Evidence from the ban of ChatGPT in Italy." *Journal of Accounting and Economics*, 80(1): 101782.

Bilinski, P. 2024. "The usefulness of ChatGPT for textual analysis of annual reports." Unpublished paper.

Blankespoor, E. 2019. "The impact of information processing costs on firm disclosure choice: Evidence from the XBRL mandate." *Journal of Accounting Research*, 57(4): 919–967.

Blankespoor, E, E DeHaan, J Wertz, and C Zhu. 2018. "Why do individual investors disregard accounting information? The roles of information awareness and acquisition costs." *Journal of Accounting Research*, 57(1): 53–84.

Blankespoor, E, G. S Miller, and H. D White. 2014. "The role of dissemination in market liquidity: Evidence from firms' use of Twitter." *Accounting Review*, 89(1): 79–112.

Bodnaruk, A. 2009. "Proximity always matters: Local bias when the set of local companies changes." *Review of Finance*, 13(4): 629–656.

Bushee, B. J, and C Leuz. 2005. "Economic consequences of SEC disclosure regulation: Evidence from the OTC bulletin board." *Journal of Accounting and Economics*, 39(2): 233–264.

Caetano, C, B Callaway, S Payne, and H. S Rodrigues. 2024. "Difference in Differences with Time-Varying Covariates." UnpublishedPaper.

Chang, Y.-C, A Ljungqvist, and K Tseng. 2023. "Do corporate disclosures constrain strategic analyst behavior?" *The Review of Financial Studies*, 36(8): 3163–3212.

Cheng, Q, P Lin, and Y Zhao. 2025. "Does Generative AI facilitate investor trading? Early evidence from ChatGPT outages." *Journal of Accounting and Economics*, 80(2–3).

Cicero, D. C, A Puckett, A. Y Wang, and S Zhang. 2024. "TAXI! Do mutual funds pursue and exploit information on local companies?" *Journal of Financial and Quantitative Analysis*, 59(7): 3340–3375.

Cohen, L. 2009. "Loyalty-based portfolio choice." *Review of Financial Studies*, 22(3): 1213–1245.

Cohen, L, C Malloy, and Q Nguyen. 2020. "Lazy prices." *Journal of Finance*, 75(3): 1371–1415.

Coval, J. D, and T. J Moskowitz. 1999. "Home bias at home: Local equity preference in domestic portfolios." *Journal of Finance*, 54(6): 2045–2073.

Coval, J. D, and T. J Moskowitz. 2001. "The geography of investment: Informed trading and asset prices." *Journal of Political Economy*, 109(4): 811–841.

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

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

Da, Z, U. G Gurun, B Li, and M Warachka. 2021. "Investment in a smaller world: The implications of air travel for investors and firms." *Management Science*, 67(1): 417–435.

Dong, Y, O. Z Li, Y Lin, and C Ni. 2016. "Does information-processing cost affect firm-specific information acquisition? Evidence from XBRL adoption." *The Journal of Financial and Quantitative Analysis*, 51(2): 435–462.

French, K. R, and J. M Poterba. 1991. "Investor diversification and international equity markets."

American Economic Review, Papers and Proceedings, 81(2): 222–226.

Gao, M, and J Huang. 2020. "Informing the market: The effect of modern information technologies on information production." *The Review of Financial Studies*, 33(4): 1367–1411.

Grinblatt, M, and M Keloharju. 2001. "How distance, language, and culture influence stockholdings and trades." *Journal of Finance*, 56(3): 1053–1073.

Gustafson, M, J. J Henry, E Kim, and K Pisciotta. 2025. "Weathering the Blow of a Disrupted Roadshow." Unpublished Paper.

Hau, H. 2001. "Location matters: An examination of trading profits." *Journal of Finance*, 56(5): 1959–1983.

He, X, A Hu, and Z Li. 2026. "Inelastic demand at the extensive and intensive margins." Unpublished paper.

Huberman, G. 2001. "Familiarity breeds investment." *Review of Financial Studies*, 14(3): 659–680.

Ivkovic, Z, and S Weisbenner. 2005. "Local does as local is: Information content of the geography of individual investors' common stock investments." *Journal of Finance*, 60(1): 267–306.

Kacperczyk, M, S Van Nieuwerburgh, and L Veldkamp. 2016. "A rational theory of mutual funds' attention allocation." *Econometrica*, 84(2): 571–626.

Kang, B. U, and J. S Nam. 2025. "Do local newspapers matter to institutional investors?" *Contemporary Accounting Research*, 42(3): 1713–1743.

Korniotis, G. M, and A Kumar. 2013. "State-level business cycles and local return predictability." *The Journal of Finance*, 68(3): 1037–1096.

Kuchler, T, Y Li, L Peng, J Stroebel, and D Zhou. 2022. "Social proximity to capital: Implications for investors and firms." *Review of Financial Studies*, 35(6): 2743–2789.

Lee, R. 2026. "Face-to-face social interactions and local information advantage." *Journal of Financial Economics, forthcoming*.

Lopez-Lira, A, and Y Tang. 2026. "Can ChatGPT forecast stock price movements? Return predictability and large language models." *Journal of Financial Economics, forthcoming*.

Loughran, T, and B McDonald. 2011. "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks." *Journal of Finance*, 66(1): 35–65.

Massa, M, and A Simonov. 2006. "Hedging, familiarity and portfolio choice." *Review of Financial*

Studies, 19(2): 633–685.

Morse, A, and S Shive. 2011. "Patriotism in your portfolio." *Journal of Financial Markets*, 14(2): 411–440.

Peress, J. 2014. "The media and the diffusion of information in financial markets: Evidence from newspaper strikes." *Journal of Finance*, 69(5): 2007–2043.

Pool, V. K, N Stoffman, and S. E Yonker. 2012. "No place like home: Familiarity in mutual fund manager portfolio choice." *Review of Financial Studies*, 25(8): 2563–2599.

Seasholes, M. S, and N Zhu. 2010. "Individual investors and local bias." *Journal of Finance*, 65(5): 1987–2010.

Shao, R, and N Wang. 2021. "Trust and local bias of individual investors." *Journal of Banking and Finance*, 133: 106273.

Sheng, J, Z Sun, B Yang, and A. L Zhang. 2026. "Generative AI and asset management." *Review of Financial Studies, forthcoming*.

Sialm, C, Z Sun, and L Zheng. 2020. "Home bias and local contagion: Evidence from funds of hedge funds." *Review of Financial Studies*, 33(10): 4771–4810.

Teo, M. 2009. "The geography of hedge funds." *Review of Financial Studies*, 22(9): 3531–3561.

Van Nieuwerburgh, S, and L Veldkamp. 2009. "Information immobility and the home bias puzzle." *Journal of Finance*, 64(3): 1187–1215.

Van Nieuwerburgh, S, and L Veldkamp. 2010. "Information acquisition and under-diversification." *Review of Economic Studies*, 77(2): 779–805.

Xu, Y, Y Xuan, and G Zheng. 2021. "Internet searching and stock price crash risk: Evidence from a quasi-natural experiment." *Journal of Financial Economics*, 141(1): 255–275.

Zeng, S, and K Zhou. 2026. "Textual analysis by hedge funds." *Journal of Accounting Research, forthcoming*.

Table 1: Sample composition

This table reports the average aggregate total net assets (TNA), the average fund TNA, and the average number of funds per quarter by Lipper investment style classification for the sample period from the second quarter of 2021 through the second quarter of 2024, with the fourth quarter of 2022 excluded as the transition quarter. Aggregate TNA and fund TNA are reported in millions of U.S. dollars. The final row reports the corresponding averages for the full sample.

Lipper Category	Sample avg. aggregate TNA per quarter ($ millions)	Sample avg. fund TNA per quarter ($ millions)	Sample avg. funds per quarter
Large-Cap Core	598,489	5,471	111
Large-Cap Growth	660,305	5,453	124
Large-Cap Value	177,792	3,320	54
Mid-Cap Core	113,812	1,527	75
Mid-Cap Growth	226,910	2,830	80
Mid-Cap Value	77,986	2,447	32
Multi-Cap Core	131,911	1,200	111
Multi-Cap Growth	219,621	2,419	91
Multi-Cap Value	270,537	1,933	141
Small-Cap Core	130,813	772	171
Small-Cap Growth	158,964	1,108	145
Small-Cap Value	22,864	532	42
Equity-Cap Income	326,051	3,623	92
Total	3,116,055	2,485	1,269

Table 2: Summary statistics of fund and firm characteristics

This table reports summary statistics for the main variables used in our analysis. Panel A reports summary statistics for variables measured at the fund–firm level. Portfolio weight is the percentage of a fund's TNA invested in a given stock and is set equal to zero in quarters in which the fund holds no position in the stock. Panel B reports summary statistics for fund-level variables. Fund size is the natural logarithm of fund TNA. Fund flow is the quarterly net flow scaled by beginning-of-quarter TNA. Fund return is the quarterly fund return. Fund expense ratio is the ratio of annual management, administrative, and 12b-1 fees and expenses to year-end TNA. Fund for institutionals is an indicator equal to one if a fund's largest share class (by TNA) is flagged as targeted to institutional investors, and zero otherwise. Panel C reports summary statistics for firm-level variables. Firm book-to-market ratio is the book value of equity divided by market value of equity. Firm return on assets is the ratio of net income to total assets. Firm debt-to-assets ratio is the ratio of the sum of short-term and long-term debt to total assets. Firm size is the natural logarithm of market capitalization. Firm prior 12-month return momentum is the stock's cumulative return over the preceding 12 months, excluding the current month, while firm quarterly return is the stock's return over the current quarter.

				Percentiles		
	#Obs	Mean	Std. Dev.	10^{th}	50^{th}	90^{th}
Panel A: Fund-firm observations						
Portfolio weight (%)	2756468	0.454	0.964	0	0	1.490
Panel B: Fund-level variables						
Fund return	14884	0.030	0.093	-0.096	0.041	0.136
Fund flow	14878	0.074	4.062	-0.078	-0.018	0.058
Fund turnover ratio	14762	0.535	0.744	0.120	0.390	1.000
Fund expense ratio	14862	0.010	0.003	0.006	0.009	0.013
Fund for institutionals	14980	0.746	0.435	0	1	1
Fund size	14985	5.935	2.107	3.277	5.966	8.627
Panel C: Firm-level variables						
Firm Book-to-Market ratio	37232	0.562	0.885	0.063	0.438	1.205
Firm Return on Assets	37348	-0.010	0.238	-0.078	0.004	0.035
Firm Debt-to-Assets ratio	37367	0.597	0.487	0.196	0.586	0.915
Firm size	37280	7.388	2.067	4.741	7.375	10.114
Firm prior 12-month return momentum	36367	0.202	1.184	-0.491	0.047	0.845
Firm quarterly return	37713	0.003	0.284	-0.276	-0.011	0.260

Table 3: The emergence of GenAI and the local bias

This table reports the results on the effect of the release of ChatGPT on mutual funds' portfolio allocations to local firms. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. Fund controls include fund return, fund flows, turnover ratio, expense ratio, institutional-fund dummy, and fund size. Firm controls include the book-to-market ratio, return on assets, debt-to-assets ratio, firm size, and prior 12-month return momentum. Columns (1)–(2) include fund, firm, and year-quarter fixed effects; columns (3)–(4) include fund-by-industry, firm, and year-quarter fixed effects; and columns (5)–(6) include fund-by-firm and year-quarter fixed effects. The sample period runs from the second quarter of 2021 through the second quarter of 2024, with the fourth quarter of 2022 excluded as the transition quarter. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.038 ***	0.038 ***	0.043 ***	0.044 ***		
	4.841	4.772	5.562	5.695		
Local × Post GPT	-0.019 ***	-0.014 ***	-0.020 ***	-0.016 ***	-0.017 ***	-0.012 ***
	-3.603	-2.613	-3.929	-2.964	-3.535	-2.581
Fund return		-0.039 ***		-0.040 ***		-0.038 ***
		-7.718		-8.197		-10.006
Fund flows		-0.000		-0.000		-0.000 **
		-1.368		-1.418		-2.081
Fund turnover ratio		-0.005 **		-0.005 **		-0.005 ***
		-2.345		-2.526		-3.161
Fund expense ratio		-0.711 **		-0.710 **		-0.522 *
		-2.007		-2.157		-1.687
Fund for institutionals		-0.002		-0.001		0.001
		-0.574		-0.477		0.458
Fund size		-0.011 ***		-0.011 ***		-0.009 ***
		-4.494		-4.598		-4.082
Firm Book-to-Market ratio		0.006 *		0.006 **		0.006 **
		1.929		1.979		2.011
Firm Return on Assets		0.172 ***		0.169 ***		0.165 ***
		6.861		6.864		7.043
Firm debt-to-assets ratio		0.074 ***		0.074 ***		0.080 ***
		7.367		7.515		8.498
Firm size		0.212 ***		0.213 ***		0.215 ***
		30.186		30.150		29.662
Firm prior 12-month return momentum		0.014 ***		0.014 ***		0.014 ***
		8.592		8.566		8.362
Fund Fixed Effects	Yes	Yes	No	No	No	No
Firm Fixed Effects	Yes	Yes	Yes	Yes	No	No
Fund-by-Industry Fixed Effects	No	No	Yes	Yes	No	No
Fund-by-Firm Fixed Effects	No	No	No	No	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	2,756,466	2,594,845	2,744,641	2,594,775	2,752,264	2,586,870
Adjusted R^2	0.209	0.214	0.263	0.266	0.776	0.781

Table 4: The emergence of GenAI and the local bias – Sample restrictions and controlling for stock returns

This table examines the robustness of the effect of the release of ChatGPT on mutual funds' portfolio allocations to local firms to alternative sample restrictions and the inclusion of contemporaneous stock returns. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. Panel A examines alternative sample restrictions. Columns (1)–(3) exclude funds located in Boston and New York, while columns (4)–(6) exclude technology firms. Panel B additionally controls for the firm's contemporaneous quarterly stock return. All specifications include fund- and firm-level controls, although for brevity these coefficient estimates are not reported. Different columns employ alternative sets of fixed effects, as indicated in the table. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A

	Excl. Boston & NY funds			Excl. Tech Firms		
	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.047 ***	0.048 ***		0.042 ***	0.045 ***	
	4.288	4.705		5.169	5.612	
Local × Post GPT	-0.021 ***	-0.022 ***	-0.022 ***	-0.015 ***	-0.015 ***	-0.012 **
	-2.774	-2.971	-3.076	-2.684	-2.797	-2.438
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	No	Yes	No	No
Firm Fixed Effects	Yes	Yes	No	Yes	Yes	No
Fund-by-Industry Fixed Effects	No	Yes	No	No	Yes	No
Fund-by-Firm Fixed Effects	No	No	Yes	No	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	1,736,933	1,736,872	1,731,215	2,194,812	2,194,742	2,188,178
Adjusted R^2	0.200	0.256	0.771	0.178	0.237	0.756

Panel B

	(1)	(2)	(3)
Local	0.038 ***	0.044 ***	
	4.805	5.734	
Local × Post GPT	-0.014 ***	-0.016 ***	-0.013 *
	-2.699	-3.060	-1.937
Firm quarterly return	0.248 ***	0.248 ***	0.249 ***
	33.270	33.350	14.690
Fund Controls	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	No
Firm Fixed Effects	Yes	Yes	No
Fund-by-Industry Fixed Effects	No	Yes	No
Fund-by-Firm Fixed Effects	No	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes
# of Observations	2,594,830	2,594,760	2,586,855
Adjusted R^2	0.216	0.268	0.784

Table 5: The emergence of GenAI and the local bias – Saturated fixed effects

This table reports the results on the effect of the release of ChatGPT on mutual funds' portfolio allocations to local firms using more saturated fixed-effect specifications. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. Column (1) includes firm, fund-by-industry, fund-by-year-quarter, and industry-by-year-quarter fixed effects. Column (2) includes fund-by-industry, fund-by-year-quarter, and firm-by-year-quarter fixed effects. Column (3) includes firm and fund-by-industry-by-year-quarter fixed effects. Firm-level controls are included in columns (1) and (3) but are absorbed by the fixed effects in column (2); fund-level controls are absorbed by the fixed effects in all specifications. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
Local	0.039 ***	0.037 ***	0.039 ***
	5.022	4.762	4.902
Local × Post GPT	-0.013 **	-0.010 *	-0.012 **
	-2.482	-1.674	-2.284
Fund Controls	No	No	No
Firm Controls	Yes	No	Yes
Firm Fixed Effects	Yes	No	Yes
Fund-by-Industry Fixed Effects	Yes	Yes	No
Fund-by-Year-Quarter Fixed Effects	Yes	Yes	No
Industry-by-Year-Quarter Fixed Effects	Yes	No	No
Firm-by-Year-Quarter Fixed Effects	No	Yes	No
Fund-by-Industry-by-Year-Quarter Fixed Effects	No	No	Yes
# of Observations	2,708,123	2,742,079	2,699,778
Adjusted R^2	0.264	0.269	0.243

Table 6: The emergence of GenAI and the local bias – Dynamic analysis

This table reports the evolution of mutual funds' portfolio allocations to local firms around the release of ChatGPT using a dynamic event-study specification. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Event time τ denotes the number of quarters relative to the release quarter of ChatGPT, $2022Q4$. The release quarter is treated as a transition period and therefore excluded from the analysis. $\tau = -1$ corresponds to $2022Q3$ and is omitted as the reference period. The coefficients on Local $\times \tau$ capture the change in the difference in portfolio weights between local and distant stocks in each event quarter relative to $2022Q3$. Columns (1) and (2) include fund, firm, and year-quarter fixed effects, while columns (3) and (4) include fund-by-firm and year-quarter fixed effects. The sample period runs from the second quarter of 2021 through the second quarter of 2024. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Local × τ = -6	0.001	-0.002	0.002	-0.002	0.001	-0.003
	0.198	-0.358	0.243	-0.281	0.105	-0.579
Local × τ = -5	0.008	0.004	0.007	0.004	0.006	0.003
	1.541	0.807	1.258	0.682	1.251	0.515
Local × τ = -4	0.002	-0.003	0.001	-0.003	0.002	-0.003
	0.468	-0.546	0.295	-0.567	0.432	-0.748
Local × τ = -3	-0.005	-0.006	-0.003	-0.005	-0.003	-0.005
	-1.262	-1.312	-0.865	-1.112	-0.864	-1.285
Local × τ = -2	0.001	0.003	0.001	0.003	0.003	0.003
	0.289	0.908	0.424	0.992	0.988	1.272
Local × τ = +1	-0.004	-0.001	-0.003	-0.001	0.000	0.003
	-0.857	-0.130	-0.774	-0.152	0.104	0.885
Local × τ = +2	-0.010 **	-0.009	-0.012 **	-0.011 **	-0.010 **	-0.009 **
	-1.981	-1.613	-2.514	-2.057	-2.252	-2.118
Local × τ = +3	-0.016 ***	-0.013 **	-0.017 ***	-0.014 ***	-0.012 ***	-0.010 **
	-2.735	-2.165	-3.059	-2.487	-2.584	-2.135
Local × τ = +4	-0.020 ***	-0.016 ***	-0.022 ***	-0.018 ***	-0.019 ***	-0.017 ***
	-3.346	-2.703	-3.856	-3.148	-3.740	-3.262
Local × τ = +5	-0.029 ***	-0.026 ***	-0.030 ***	-0.027 ***	-0.028 ***	-0.025 ***
	-4.567	-3.919	-4.881	-4.359	-4.898	-4.495
Local × τ = +6	-0.029 ***	-0.025 ***	-0.031 ***	-0.026 ***	-0.028 ***	-0.025 ***
	-4.164	-3.463	-4.579	-3.763	-4.393	-3.928
Fund Controls	No	Yes	No	Yes	No	Yes
Firm Controls	No	Yes	No	Yes	No	Yes
Fund Fixed Effects	Yes	Yes	No	No	No	No
Firm Fixed Effects	Yes	Yes	Yes	Yes	No	No
Fund-by-Industry Fixed Effects	No	No	Yes	Yes	No	No
Fund-by-Firm Fixed Effects	No	No	No	No	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	2,756,466	2,594,845	2,744,641	2,594,775	2,752,264	2,586,870
Adjusted R^2	0.209	0.214	0.263	0.266	0.776	0.781

Table 7: The emergence of GenAI and the local bias – Extensive vs. intensive margin

This table examines whether the relative reduction in mutual funds' local portfolio allocations following the release of ChatGPT operates through the extensive or intensive margin. In column (1), the dependent variable is an indicator equal to one if the fund holds a given stock in the current quarter and zero otherwise. In column (2), the sample is restricted to fund–firm pairs that were not held in the previous quarter and are therefore at risk of entry; the dependent variable is an indicator equal to one if the fund initiates a position in the current quarter. In column (3), the sample is restricted to fund–firm pairs that were held in the previous quarter and are therefore at risk of exit; the dependent variable is an indicator equal to one if the fund fully liquidates the position in the current quarter. In column (4), the sample is restricted to fund–firm pairs held in consecutive quarters, and the dependent variable is the change in portfolio weight. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in the respective holding decision or portfolio allocation for local firms relative to distant firms following the release of ChatGPT. All specifications include fund- and firm-level controls, fund and firm fixed effects, and year-quarter fixed effects. For brevity, coefficients on the control variables are not reported. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1) Extensive Margin	(2) Extensive Margin In	(3) Extensive Margin Out	(4) Intensive Margin
Local	0.023 ***	0.005 ***	-0.002	0.006 ***
	6.043	3.298	-1.047	3.213
Local × Post GPT	-0.009 **	-0.004 *	0.000	-0.012 ***
	-2.071	-1.928	0.086	-4.018
Fund Controls	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes
# of Observations	2,594,845	1,204,790	872,588	781,157
Adjusted R^2	0.126	0.043	0.096	0.030

Table 8: The emergence of GenAI and the local bias – Firm complexity and size

This table shows whether the post-ChatGPT decline in mutual funds' local overweighting varies with firms' disclosure complexity. Firm disclosure complexity are measured using filing size and the SMOG readability index constructed from firms' 10-K filings. Filing size captures the volume of information investors must process, while a higher SMOG score indicates that greater reading effort is required to understand the disclosure. In Panel A, the sample is split into low- and high-processing-cost groups based on the median of the pre-ChatGPT average of each complexity measure, and the baseline specification is re-estimated separately for each group. Panel B further sorts firms independently by disclosure complexity and firm size and re-estimates the baseline specification separately for each of the resulting four size–complexity groups. For both filing size and SMOG, Low and High refer to disclosure complexity, while Small and Large refer to firm size. In all specifications, the dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. All specifications include fund- and firm-level controls, fund fixed effects, firm fixed effects, and year-quarter fixed effects. For brevity, coefficients on the control variables are not reported. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Split by firm complexity

	File size		SMOG	
	(1)	(2)	(3)	(4)
	Low	High	Low	High
Local	0.027 **	0.042 ***	0.029 **	0.040 ***
	2.055	4.698	2.540	4.133
Local × Post GPT	0.005	-0.019 ***	0.004	-0.022 ***
	0.559	-3.153	0.513	-3.341
Fund Controls	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes
# of Observations	849,551	1,711,603	819,284	1,741,992
Adjusted R^2	0.256	0.199	0.282	0.190

Panel B: Split by firm complexity and size

	File size				SMOG			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Low / Small	Low / Large	High / Small	High / Large	Low / Small	Low / Large	High / Small	High / Large
Local	-0.002	0.034 **	0.003	0.044 ***	-0.003	0.035 ***	0.002	0.042 ***
	-0.122	2.245	0.142	4.713	-0.186	2.683	0.074	4.185
Local × Post GPT	0.016	0.003	0.006	-0.021 ***	0.021	0.002	-0.004	-0.023 ***
	0.945	0.251	0.287	-3.263	1.455	0.156	-0.186	-3.397
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	137,170	712,170	108,459	1,603,141	144,641	674,534	101,111	1,640,777
Adjusted R^2	0.369	0.251	0.364	0.196	0.353	0.280	0.387	0.187

Table 9: The emergence of GenAI and the local bias – Quantitative vs. discretionary funds

This table shows whether the post-ChatGPT decline in mutual funds' local overweighting differs between quantitative and discretionary funds. Building on Abis (2022), funds are classified as quantitative or discretionary using ChatGPT 5.5 to analyze the Principal Investment Strategy (PIS) sections of fund prospectuses. Column (1) reports coefficient estimates for quantitative funds, for which statistical, algorithmic, or other model-based methods determine security selection and portfolio construction, while column (2) reports coefficient estimates for discretionary funds, for which human judgment is indispensable to investment decisions. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. Both specifications include fund- and firm-level controls, fund fixed effects, firm fixed effects, and year-quarter fixed effects. For brevity, coefficients on the control variables are not reported. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1) Quantitative Funds	(2) Discretionary Funds
Local	0.010	0.040 ***
	0.626	4.466
Local × Post GPT	-0.006	-0.013 **
	-0.681	-2.171
Fund Controls	Yes	Yes
Firm Controls	Yes	Yes
Fund Fixed Effects	Yes	Yes
Firm Fixed Effects	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes
# of Observations	245,154	2,246,513
Adjusted R^2	0.268	0.216

Table 10: The emergence of GenAI and the local bias – Fund managerial structure

This table shows whether the post-ChatGPT decline in mutual funds' local overweighting varies with fund managerial structure. The fund sample is split into team-managed and non-team-managed funds based on fund manager information reported in CRSP. Column (1) reports coefficient estimates for team-managed funds, while column (2) reports coefficient estimates for non-team-managed funds. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. Both specifications include fund- and firm-level controls, fund fixed effects, firm fixed effects, and year-quarter fixed effects. For brevity, coefficients on the control variables are not reported. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1) Team-managed Funds	(2) Non-team-managed Funds
Local	0.014	0.048 ***
	1.207	4.630
Local × Post GPT	-0.008	-0.019 ***
	-1.136	-2.624
Fund Controls	Yes	Yes
Firm Controls	Yes	Yes
Fund Fixed Effects	Yes	Yes
Firm Fixed Effects	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes
# of Observations	553,561	1,894,721
Adjusted R^2	0.231	0.225

Table 11: The emergence of GenAI and the stock-picking ability of local vs. distant holdings

This table reports the results on the effect of the release of ChatGPT on mutual funds' stock-picking ability in local and distant portfolio segments. The dependent variable is stock-picking ability, calculated separately for each fund's local and distant holdings following Kacperczyk, Van Nieuwerburgh and Veldkamp (2016). Stock-picking ability is higher when a fund actively overweights stocks that subsequently realize high idiosyncratic returns and underweights stocks that subsequently realize low idiosyncratic returns. Local is an indicator equal to one for the local portfolio segment and zero for the distant segment. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in stock-picking ability of local portfolios relative to distant portfolios following the release of ChatGPT. Column (1) includes fund and year-quarter fixed effects without fund-level controls. Column (2) adds fund-level controls, including fund flows, turnover ratio, expense ratio, institutional-fund dummy, and fund size. Column (3) includes fund-level controls and investment-style-by-year-quarter fixed effects. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
Local	0.0001 ***	0.0001 ***	0.0001 ***
	3.6263	3.3836	3.5884
Local × Post GPT	-0.0002 ***	-0.0002 ***	-0.0002 ***
	-3.7381	-3.3453	-3.5881
Fund flows		-0.0000 ***	0.0000
		-5.6441	1.0855
Fund turnover ratio		0.0001	-0.0000 *
		1.2090	-1.6724
Fund expense ratio		0.0076	-0.0032
		0.6281	-0.8379
Fund for institutionals		-0.0000	-0.0001 ***
		-0.1509	-2.7858
Fund size		-0.0006 ***	0.0000
		-6.7646	0.2822
Fund Fixed Effects	Yes	Yes	No
Year-Quarter Fixed Effects	Yes	Yes	No
Style-by-Year-Quarter Fixed Effects	No	No	Yes
# of Observations	26,567	25,575	25,584
Adjusted R^2	0.2873	0.2894	0.3319

Table 12: The emergence of GenAI and local vs. distant portfolio performance

This table reports the results on the effect of the release of ChatGPT on the performance of mutual funds' local and distant portfolio holdings. The dependent variable is the return on the fund's local or distant portfolio segment, constructed from end-of-quarter holdings. Columns (1)–(3) use portfolio excess returns over the risk-free rate, while columns (4)–(6) use DGTW-adjusted returns. Local is an indicator equal to one for the local portfolio segment and zero for the distant segment. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures how the performance gap between local and distant portfolio holdings changes following the release of ChatGPT. Columns (1) and (4) include fund and year-quarter fixed effects without fund-level controls. Columns (2) and (5) additionally include fund-level controls, including fund flows, turnover ratio, expense ratio, institutional-fund dummy, and fund size. Columns (3) and (6) include fund-level controls and investment-style-by-year-quarter fixed effects. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Excess return			DGTW-adjusted return		
	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.0052 ***	0.0053 ***	0.0055 ***	-0.0009	-0.0007	-0.0005
	4.8549	4.8451	5.0569	-0.9139	-0.6651	-0.5114
Local × Post GPT	-0.0160 ***	-0.0156 ***	-0.0157 ***	-0.0032 **	-0.0031 **	-0.0033 **
	-9.6009	-9.1331	-9.2394	-2.1784	-2.0458	-2.1634
Fund flows		-0.0005 ***	-0.0000		-0.0002 ***	0.0000
		-8.6199	-0.7287		-4.2566	1.0384
Fund turnover ratio		-0.0008	-0.0010 **		0.0016	-0.0002
		-0.4055	-2.0542		1.3960	-0.5277
Fund expense ratio		0.1423	-0.2420		0.2969	-0.1071
		0.3105	-1.6391		0.8594	-0.8436
Fund for institutionals		-0.0035	-0.0021 **		-0.0003	-0.0016 **
		-0.7333	-2.2749		-0.0910	-1.9753
Fund size		-0.0141 ***	0.0001		-0.0098 ***	-0.0002
		-6.9035	0.3170		-6.5804	-1.1380
Fund Fixed Effects	Yes	Yes	No	Yes	Yes	No
Year-Quarter Fixed Effects	Yes	Yes	No	Yes	Yes	No
Style-by-Year-Quarter Fixed Effects	No	No	Yes	No	No	Yes
# of Observations	28,161	27,091	27,093	28,161	27,091	27,093
Adjusted R^2	0.2829	0.2840	0.3412	0.0243	0.0260	0.0608

Table 13: The emergence of GenAI and the local bias – Wider "local" definitions

This table examines the robustness of the effect of the release of ChatGPT on mutual funds' local overweighting to wider definitions of geographic proximity. Columns (1)–(3) define a firm as local if it is located within 150 miles of the fund advisor, while columns (4)–(6) use a 250-mile threshold. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters falls within the specified distance threshold, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. All specifications include fund- and firm-level controls, although for brevity these coefficient estimates are not reported. Different columns employ alternative sets of fixed effects, as indicated in the table. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Local is within 150 miles			Local is within 250 miles		
	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.036 ***	0.040 ***		0.027 ***	0.029 ***	
	4.858	5.649		4.391	4.905	
Local × Post GPT	-0.015 ***	-0.016 ***	-0.011 ***	-0.018 ***	-0.019 ***	-0.014 ***
	-2.984	-3.248	-2.592	-4.554	-4.887	-4.149
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	No	Yes	No	No
Firm Fixed Effects	Yes	Yes	No	Yes	Yes	No
Fund-by-Industry Fixed Effects	No	Yes	No	No	Yes	No
Fund-by-Firm Fixed Effects	No	No	Yes	No	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	2,594,845	2,594,775	2,586,870	2,594,845	2,594,775	2,586,870
Adjusted R^2	0.214	0.266	0.781	0.214	0.266	0.781

Table 14: The emergence of GenAI and the local bias – Using number of shares

This table examines whether the effect of the release of ChatGPT on mutual funds' local holdings is robust to measuring holdings using the number of shares rather than portfolio weights. The dependent variable is the logarithm of the number of shares of a given stock held by the fund. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' holdings of local stocks relative to distant stocks following the release of ChatGPT. Columns (1)–(2) include fund, firm, and year-quarter fixed effects; columns (3)–(4) include fund-by-industry, firm, and year-quarter fixed effects; and columns (5)–(6) include fund-by-firm and year-quarter fixed effects. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.264 ***	0.275 ***	0.290 ***	0.296 ***		
	6.097	6.228	6.906	6.877		
Local × Post GPT	-0.112 **	-0.091 **	-0.116 ***	-0.093 **	-0.118 ***	-0.093 **
	-2.563	-2.055	-2.681	-2.120	-2.678	-2.132
Fund Controls	No	Yes	No	Yes	No	Yes
Firm Controls	No	Yes	No	Yes	No	Yes
Fund Fixed Effects	Yes	Yes	No	No	No	No
Firm Fixed Effects	Yes	Yes	Yes	Yes	No	No
Fund-by-Industry Fixed Effects	No	No	Yes	Yes	No	No
Fund-by-Firm Fixed Effects	No	No	No	No	Yes	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	2,756,173	2,594,558	2,744,350	2,594,488	2,751,971	2,586,583
Adjusted R^2	0.136	0.145	0.168	0.177	0.538	0.555

Table 15: The emergence of GenAI and the local bias – Alternative event window

This table examines the robustness of the effect of the release of ChatGPT on mutual funds' local overweighting to alternative event windows. Columns (1)–(3) use a winder event window spanning eight quarters before to eight quarters after the release quarter, while columns (4)–(6) use a narrower window spanning four quarters before to four quarters after the release quarter. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. All specifications include fund- and firm-level controls. Different columns use alternative fixed-effect specifications, as indicated in the table. For brevity, coefficients on the control variables are not reported. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Event window [-8Q, +8Q]			Event window [-4Q, +4Q]		
	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.040 ***	0.044 ***		0.045 ***	0.054 ***	
	5.529	6.316		4.975	6.310	
Local × Post GPT	-0.015 ***	-0.016 ***	-0.014 **	-0.013 **	-0.015 ***	-0.010
	-2.912	-3.353	-2.066	-2.216	-2.681	-1.359
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	No	Yes	No	No
Firm Fixed Effects	Yes	Yes	No	Yes	Yes	No
Fund-by-Industry Fixed Effects	No	Yes	No	No	Yes	No
Fund-by-Firm Fixed Effects	No	No	Yes	No	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	3,765,146	3,764,771	3,758,324	1,538,479	1,538,446	1,534,269
Adjusted R^2	0.205	0.252	0.753	0.225	0.284	0.815

Table 16: The emergence of GenAI and the local bias – Alternative clustering specifications

This table examines the robustness of the effect of the release of ChatGPT on mutual funds' local overweighting to alternative clustering specifications. Columns (1)–(3) cluster standard errors at the stock level, while columns (4)–(6) use two-way clustering at the fund and stock levels. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. All specifications include fund- and firm-level controls. Different columns employ alternative sets of fixed effects, as indicated in the table. For brevity, coefficients on the control variables are not reported. *t*-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Stock-level Clustering			Fund- & Stock-level Clustring		
	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.038 ***	0.044 ***		0.038 ***	0.044 ***	
	4.036	4.401		3.659	4.089	
Local × Post GPT	-0.014 **	-0.016 **	-0.012 *	-0.014 *	-0.016 **	-0.012 *
	-2.030	-2.209	-1.835	-1.911	-2.110	-1.804
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	No	Yes	No	No
Firm Fixed Effects	Yes	Yes	No	Yes	Yes	No
Fund-by-Industry Fixed Effects	No	Yes	No	No	Yes	No
Fund-by-Firm Fixed Effects	No	No	Yes	No	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	2,594,845	2,594,775	2,586,870	2,594,845	2,594,775	2,586,870
Adjusted R^2	0.214	0.266	0.781	0.214	0.266	0.781

Table 17: The emergence of GenAI and the local bias – Controlling for pre-event covariates

This table examines the robustness of the effect of the release of ChatGPT on mutual funds' local overweighting to controlling for pre-event covariates. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Local × Post GPT captures the change in funds' portfolio allocation to local firms relative to distant firms following the release of ChatGPT. Fund- and firm-level controls are measured using their pre-ChatGPT values and interacted with Post GPT indicator. Different columns use alternative fixed-effect specifications, as indicated in the table. For brevity, coefficients on the control variables are not reported. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
Local	0.039 ***	0.044 ***	
	4.726	5.578	
Local × Post GPT	-0.015 ***	-0.016 ***	-0.012 **
	-2.720	-2.942	-2.409
Fund Controls	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	No
Firm Fixed Effects	Yes	Yes	No
Fund-by-Industry Fixed Effects	No	Yes	No
Fund-by-Firm Fixed Effects	No	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes
# of Observations	2,452,856	2,450,822	2,452,315
Adjusted R^2	0.210	0.261	0.771

Table 18: The emergence of GenAI and the local bias – Placebo "local" definition

This table reports placebo tests of the effect of the release of ChatGPT on mutual funds' local overweighting using alternative placebo definitions of local fund–firm pairs. Columns (1)–(3) define placebo local pairs as those separated by 900–1,000 miles, while columns (4)–(6) use a distance band of 1,400–1,500 miles. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Placebo Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters falls within the specified placebo distance band, and zero otherwise. Post GPT is an indicator equal to one for quarters following the release of ChatGPT. The coefficient on Placebo Local × Post GPT captures the change in funds' portfolio allocation to firms within the placebo distance band relative to firms outside that band following the release of ChatGPT. All specifications include fund- and firm-level controls, although for brevity these coefficient estimates are not reported. Different columns employ alternative sets of fixed effects, as indicated in the table. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Local is 900 to 1,000 miles			Local is 1,400 to 1,500 miles		
	(1)	(2)	(3)	(4)	(5)	(6)
Local	-0.001	0.008		0.005	0.002	
	-0.095	0.974		0.446	0.179	
Placebo Local × Post GPT	-0.003	-0.003	-0.007	0.010	0.009	0.008
	-0.475	-0.388	-1.281	1.167	1.092	1.052
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	No	Yes	No	No
Firm Fixed Effects	Yes	Yes	No	Yes	Yes	No
Fund-by-Industry Fixed Effects	No	Yes	No	No	Yes	No
Fund-by-Firm Fixed Effects	No	No	Yes	No	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	2,594,845	2,594,775	2,586,870	2,594,845	2,594,775	2,586,870
Adjusted R^2	0.214	0.266	0.781	0.214	0.266	0.781

Table 19: The emergence of GenAI and the local bias – Placebo event dates

This table reports placebo tests of the effect of the release of ChatGPT on mutual funds' local overweighting using alternative pseudo-event dates prior to the actual release of ChatGPT. Columns (1)–(2) set the pseudo-event date to July 1, 2018; columns (3)–(4) to January 1, 2019; and columns (5)–(6) to July 1, 2019. The dependent variable is the portfolio weight, measured as the percentage of a fund's total net assets invested in a given stock. Local is an indicator equal to one if the distance between the fund advisor and the firm's headquarters is within 100 miles, and zero otherwise. Placebo Post is an indicator equal to one for quarters following the corresponding pseudo-event date. The coefficient on Local × Placebo Post captures the change in funds' portfolio allocation to local firms relative to distant firms following each pseudo-event date. All specifications include fund- and firm-level controls, although for brevity these coefficient estimates are not reported. Different columns employ alternative sets of fixed effects, as indicated in the table. Standard errors are clustered at the fund level, and t-statistics are reported below the coefficient estimates. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	Pseudo date 01/07/2018		Pseudo date 01/01/2019		Pseudo date 01/07/2019	
	(1)	(2)	(3)	(4)	(5)	(6)
Local	0.028 ***		0.027 ***		0.031 ***	
	4.206		4.092		4.709	
Local × Placebo Post	0.004	0.003	0.002	0.003	-0.005	-0.003
	0.887	0.353	0.540	0.333	-1.140	-0.406
Fund Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fund Fixed Effects	Yes	No	Yes	No	Yes	No
Firm Fixed Effects	Yes	No	Yes	No	Yes	No
Fund-by-Industry Fixed Effects	No	No	No	No	No	No
Fund-by-Firm Fixed Effects	No	Yes	No	Yes	No	Yes
Year-Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
# of Observations	2,832,378	2,821,293	2,875,617	2,859,634	2,795,863	2,784,129
Adjusted R^2	0.175	0.738	0.177	0.740	0.183	0.747

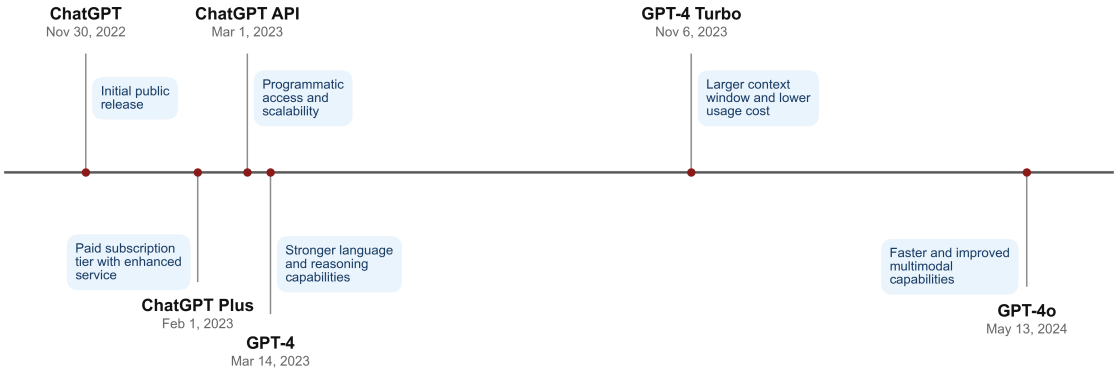


Figure 1: Timeline of ChatGPT and related OpenAI models. This figure summarizes major milestones in the development of ChatGPT and related OpenAI models over the sample period. The timeline highlights the initial release of ChatGPT, the introduction of ChatGPT Plus and API access, and subsequent releases of GPT-4, GPT-4 Turbo, and GPT-4o, which progressively expanded model capability, availability, and scalability. Although several competing large language models—including Google's Bard (later Gemini), Anthropic's Claude, Microsoft's Bing Chat (later Copilot), and Meta's LLaMA—entered the market during the sample period, ChatGPT-based tools remained the dominant platform in terms of adoption. Our empirical analysis focuses on the ChatGPT-3.5 and GPT-4 periods.

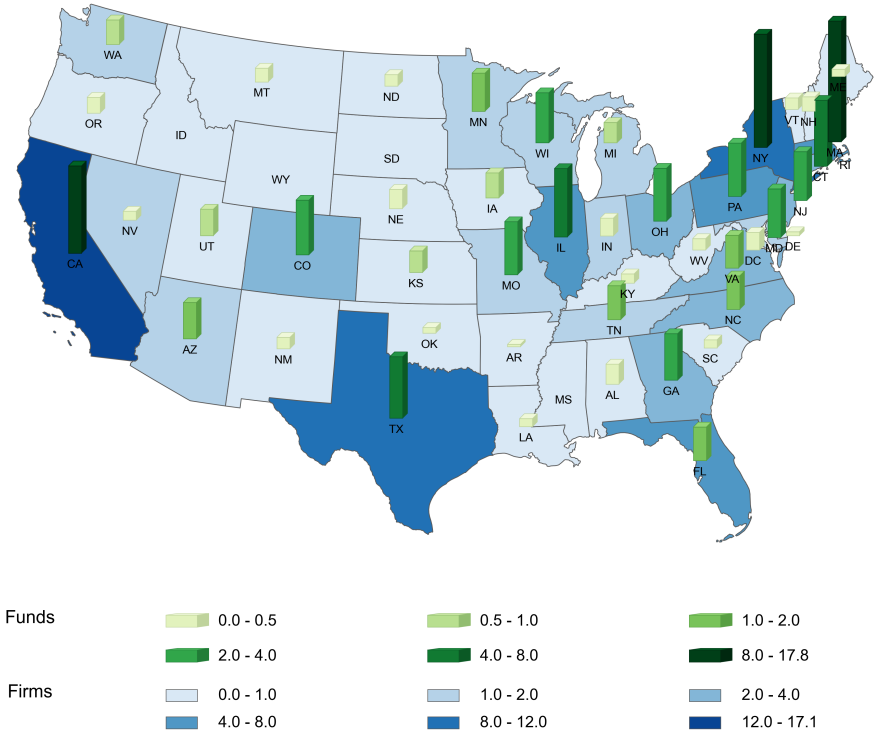


Figure 2: Geographic distribution of U.S. active equity mutual fund advisors and portfolio firms. The figure shows the time-series average share of fund advisors located in each state (indicated by bars) and the average share of firms headquartered in each state (indicated by state shading). State-level shares are calculated quarterly and then averaged over the sample period from the second quarter of 2021 through the second quarter of 2024. Alaska, Hawaii, Puerto Rico, and the U.S. Virgin Islands are excluded. A value of zero indicates that no fund advisor or firm is located in the state during the sample period.

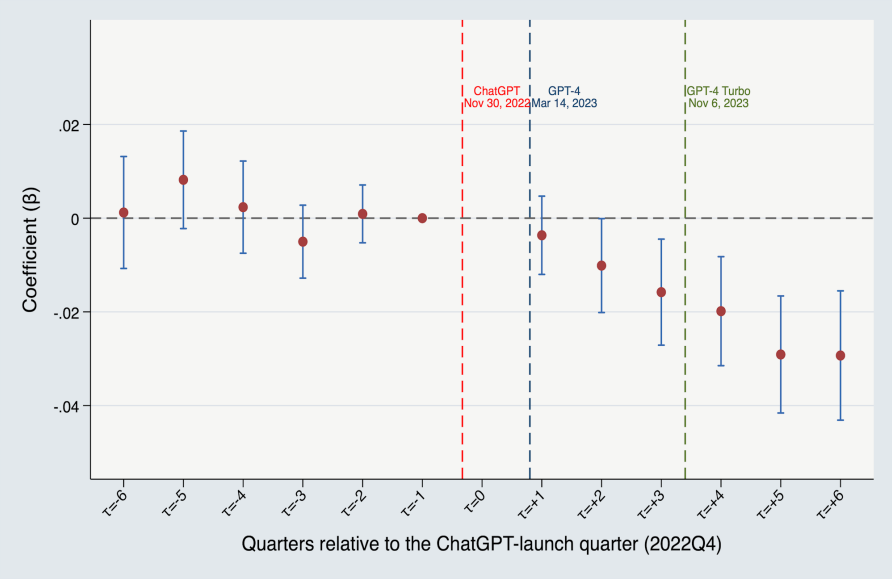


Figure 3: Evolution of local overweighting around the release of ChatGPT. This figure plots the event-study coefficients on Local $\times$ τ by estimating Equation 3, where Local equals one for fund–firm pairs located within 100 miles and τ indexes quarters relative to the ChatGPT release quarter, 2022Q4. The release quarter is excluded, and τ=$-$1 (2022Q3) is the omitted reference period. The coefficients measure changes in the difference in portfolio weights between local and distant stocks relative to 2022Q3. Vertical bars show 95% confidence intervals. The dashed vertical lines mark the releases of ChatGPT, GPT-4, and GPT-4 Turbo.