AI-Powered Analysts *

Michael Kimbrough
Robert H. Smith School of Business
University of Maryland
mkimbrou@umd.edu

Leonard Yang Liu
Robert H. Smith School of Business
University of Maryland
yliu337@umd.edu

Musa Subasi
Robert H. Smith School of Business
University of Maryland
msubasi@umd.edu

July 2026

* We appreciate the helpful comments of Sean Cao, Rebecca Hann, Mark Zakota, Nick Seybert, Xinyan Yan, Edna Lopez Avila, workshop participants at the University of Maryland, and participants at the 2024 Southern Finance Association (SFA) Annual Meeting. We are grateful to the SFA for awarding our paper the Best Paper Award in FinTech at the 2024 Annual Meeting. We gratefully acknowledge financial support from the University of Maryland for data acquisition.

AI-Powered Analysts

Abstract

We examine how brokerages' institutional adoption of artificial intelligence (AI), measured from AI skill requirements in research-related job postings, shapes sell-side analyst work. AI adoption varies across brokerages and is predicted by peer adoption, organizational capacity, analyst human capital, and institutional demand. Comparing the same analyst-firm pair over time, analysts at higher-adoption brokerages issue more accurate earnings forecasts. The gains are larger when forecasts rely more on public, machine-readable information—that is, for forecasts issued early in the period, for firms with lower information asymmetry and more readable disclosures, and by analysts with technical backgrounds. Using intraday decision fatigue, we show that AI substitutes for analysts' cognitive effort within the forecasting task — it attenuates fatigue-related accuracy losses and herding. Freed cognitive capacity is redeployed toward relationship building with management and investing clients. Forecasts by analysts at higher-adoption brokerages elicit stronger price reactions, and adopters shift hiring toward senior analysts. Institutional AI thus benefits analyst research by conserving the scarce cognition on which it runs, with gains that surface in forecast quality, market pricing, and the composition of the analyst workforce.

Keywords: artificial intelligence; analysts; financial intermediary; forecast accuracy; financial information processing; decision fatigue.

JEL Classification: G10; G14; G24; G41; O33.

Data Availability: Data used in this study are available from sources identified in the manuscript.

I. INTRODUCTION

Sell-side equity analysts do information-processing work under time pressure. They turn public filings, prices, and disclosures into the earnings forecasts that guide investors and discipline managers (e.g., Derrien and Kecskés 2013; Chen, Harford, and Lin 2015), and they do so on deadline, across many firms at once. A growing body of evidence shows that artificial intelligence (AI) is now capable at exactly this kind of work. AI models can rival or beat analysts in forecasting horseraces (van Binsbergen, Han, and Lopez-Lira 2023; Cao, Jiang, Wang, and Yang 2024), and the technology is no longer only a benchmark but a competitor: robo-analysts already issue recommendations with economically meaningful investment value (Coleman, Merkley, and Pacelli 2022), and fintech research competes with and disciplines sell-side output (Jame, Markov, and Wolfe 2022). Brokerages, in turn, are investing in the technology. Whether that investment is realized as better analyst output is not obvious. It requires well-designed tools, workflows that route the right work to the machine, and analysts willing to rely on the machine's output rather than discount it due to algorithm aversion (Dietvorst, Simmons, and Massey 2015). A capability demonstrated in forecasting horseraces need not translate to the way analysts actually produce forecasts. In this paper, we shift attention from demonstrated capability to adoption and realized output. Specifically, we examine whether a brokerage's AI adoption improves its analysts' relative forecast accuracy and, if it does, how the gain arises.

We measure a brokerage's AI adoption from its hiring. Following Babina, Fedyk, He, and Hodson (2024), we use the share of a brokerage's research-related job postings that require an AI skill, constructed from Lightcast vacancy data, as a brokerage-quarter proxy for in-house AI investment. The measure is deliberately conservative. It captures AI capability a brokerage builds internally through the roles that support analyst research, rather than the specific tools any individual analyst uses. Because brokerages that hire AI talent are also likely making broader investments in AI that our measure does not capture, our estimates likely understate the full effect of institutional AI adoption.[1] Because the proxy is indirect, we validate it by showing that

[1] Brokerages can also adopt AI through outsourcing or off-the-shelf services, commitments our hiring measure

brokerages that hire more AI talent file more AI patents, and more highly cited ones, and devote a larger share of their own earnings call discussion to AI.

We then examine what drives adoption. Organizing the determinants into the competitive environment, brokerage capacity, analyst human capital, and the demand from a brokerage's coverage portfolio, we find that adoption rises with peer adoption, internal capacity, analyst experience, and institutional demand. The demand-side drivers do not predict non-AI technology hiring, which speaks to the construct's specificity. These patterns place the measure within the forces that pull AI onto Wall Street, where competition has long shaped research production (Hong and Kacperczyk 2010; Merkley, Michaely, and Pacelli 2017; Jame et al. 2022).

Analysts employed by brokerages with greater AI adoption issue more accurate forecasts. In specifications that compare the same analyst covering the same firm over time, moving brokerage AI adoption from the sample minimum to the sample maximum is associated with approximately a 2.5 percent improvement in relative forecast accuracy. The magnitude is economically meaningful. It is roughly one-half of the improvement associated with a comparable increase in the analyst's firm-specific experience. The association is robust to a staged set of fixed effects that culminates in analyst-firm fixed effects, to longer-window and cumulative-stock versions of the AI-adoption measure, and to an absolute forecast-error specification with firm-year fixed effects.[2] We do not interpret these estimates as arising from clean exogenous variation in AI adoption. Instead, we provide complementary evidence from a shock to analysts' access to in-house AI resources. During pandemic work-from-home periods, the accuracy benefit is concentrated among analysts whose brokerage headquarters were not subject to stay-at-home orders. This pattern is consistent with the benefit operating through access to institutional AI resources rather than through coincident analyst-, firm-, or brokerage-level trends.

does not capture. To the extent such brokerages populate the comparison group, our estimates are a lower bound. Consistent with reading the measure as institutional information-processing capability, Huang, Hugon, Zhang, and Zheng (2026) find that before generative AI, the use of AI in analyst reports was negligible even as brokerages invested in AI, indicating that in-house AI served information processing rather than report writing.

[2]The alternative-measure estimates appear in Table 5, Panels B and C. The absolute-forecast-error specification and further robustness tests, including restricting to continuously active brokerages, excluding likely off-the-shelf AI adopters from the comparison group, controlling for big-data hiring, and sample splits by brokerage coverage, are reported in the Online Appendix (Section OA II).

To interpret these gains, we build on the task-based view of automation and human-AI collaboration, which develops from Autor (2015) through Acemoglu and Restrepo (2018, 2019) and Acemoglu (2025) to Fügener, Walzner, and Gupta (2026). In this view, a job is a bundle of tasks, and AI can substitute for human effort in some subtasks while complementing human effort in others. Fügener et al. (2026) distinguish two complementarities that organize our mechanism tests. The first operates within a task - AI performs some subtasks while augmenting the human judgment applied to the remaining subtasks. Forecast accuracy may therefore improve either because AI substitutes for analysts' cognitive effort or because it enhances the judgment analysts continue to supply. The second operates between tasks - capacity freed by automation is redeployed toward other activities in which analysts retain a comparative advantage.

Decision fatigue helps distinguish substitution from augmentation within the forecasting task. As analysts issue more forecasts over a trading day, accumulated depletion of cognitive resources degrades later forecasts (Hirshleifer, Levi, Lourie, and Teoh 2019; Jiao 2024). If AI primarily substitutes for analysts' depleted cognitive effort, its benefit should be largest at the fatigued end of the day. If AI primarily augments judgment-intensive inputs, its benefit should be weaker when fatigue erodes the human judgment on which such augmentation depends. The evidence is more consistent with substitution within the forecasting task. Brokerage AI offsets approximately three-fifths of the fatigue-related decline in forecast accuracy and reduces the fatigue-induced increase in herding, a low-effort heuristic, by roughly two-thirds.

We then examine whether the capacity freed by this substitution is redeployed toward activities less amenable to automation. Consistent with between-task complementarity, analysts at more AI-intensive brokerages invest more in relationship building—the relationship-based work of cultivating access to firm management and standing with the institutional clients their research serves, tasks that turn on the analyst's own preparation, judgment, and engagement rather than on information processing. Earnings conference calls offer a window into this effort because engagement on a call is a single activity with two audiences: management and clients. Conference call participation allows the analyst to deepen her relationship with management, which controls participation and the order of the Q&A queue (Mayew 2008; Cohen, Lou, and Malloy 2020), while demonstrating her preparation to the clients listening in. We find that analysts at more AI-intensive brokerages are more likely to participate in these calls, are called on earlier in the question-and-answer queue, and elicit longer responses from management.[3]

Our cross-sectional tests suggest that AI's advantage lies in processing public, machine-readable information. The accuracy benefit is somewhat larger for forecasts issued early in the cycle, when analysts rely on historical data; for firms with lower information asymmetry across analysts and with more readable disclosures, which are easier to process mechanically; and for analysts with technical backgrounds, who are less prone to algorithm aversion (Dietvorst et al. 2015).[4]

The consequences of brokerage AI adoption extend to capital and labor markets. In capital markets, forecasts from AI-powered analysts are more informative as seen in more dramatic price movements on their forecast days. This effect is concentrated among brokerages with greater AI disclosure. In labor markets, AI reshapes whom brokerages hire. Brokerages that adopt AI tilt the experience composition of their analyst hiring from entry-level toward more senior roles without reducing total hiring. The pattern is a recomposition of an expanding workforce toward judgment-intensive roles, not displacement of analysts.

Our study makes three contributions. First, we develop a validated, research-specific measure of brokerage AI adoption and provide a structured account of its determinants. Existing AI measures in finance generally capture firm-wide or bank-wide AI investment. We isolate AI adoption in roles that support analyst research, validate the proxy against patenting and disclosure,

[3]To complement our archival evidence with field evidence, we conducted supplementary interviews with sell-side analysts. Participants' descriptions of how in-house AI has changed their work are consistent with the mechanism we document: AI absorbs routine, data-intensive processing within the forecasting task, and the capacity it frees is redirected toward the relationship-based work analysts conduct with management and clients. The interview study was reviewed and approved by the University of Maryland, College Park Institutional Review Board under expedited review as minimal-risk research (Application #2698-1). Concurrent field studies reach similar takeaways: analysts surveyed and interviewed by Christ, Kim, and Yip (2025) report using AI primarily to streamline routine tasks such as text summarization and data collection, and the analysts interviewed by Shanthikumar and Yoo (2026) describe automation freeing capacity that they redirect toward information acquisition and management access.

[4]The benefit is also larger for less-experienced analysts, a pattern we report among the additional tests in the Online Appendix (Section OA I) and interpret with caution: experience is highly correlated with age, and younger analysts are likelier to use AI tools, so the pattern may reflect differences in AI usage across experience levels rather than AI leveling differences in skill. The same appendix section also reports tests of forecast characteristics, forecast optimism, alternative fixed-effects structures, and a brokerage-level version of the accuracy test.

and show that competition, capacity, and institutional demand predict the diffusion of AI into sell-side research (Hong and Kacperczyk 2010; Merkley et al. 2017; Jame et al. 2022). This measure allows us to study organizational AI resources available to analysts rather than public AI tools or machine-generated research products outside the brokerage research process. Prior work documents that AI-based competitors already operate on Wall Street (Coleman et al. 2022; Jame et al. 2022); we document how traditional analysts, and the brokerages that employ them, respond to and benefit from the technology. Collectively, these studies enhance our understanding of how AI benefits investors.

Second, we contribute to research on human-AI collaboration by identifying how AI changes the production of earnings forecasts. Existing studies often benchmark machines against analysts (van Binsbergen et al. 2023; Cao et al. 2024); we study analysts working inside institutions that have adopted AI. To our knowledge, we are the first to use decision fatigue to test whether the within-forecast AI effect reflects substitution for analysts' cognitive effort or augmentation of analysts' judgment. Because fatigue impairs the human input required for judgment augmentation, the setting separates the two. The evidence points to substitution within the forecasting task and to reallocation across tasks; AI attenuates fatigue-related accuracy losses and herding, while analysts at more AI-intensive brokerages invest more in relationship building. The closest concurrent work is Shanthikumar and Yoo (2026), who document a related automation-then-reallocation pattern in analyst research and tie its automation step to forecast timeliness. We differ by opening the forecasting task itself and by measuring the reallocation margin through Q&A queue position and management response length, rather than participation alone. The two studies are independent and complementary. That a similar pattern emerges across different samples, adoption measures, and research designs strengthens the evidence that institutional AI is reshaping analyst work, beyond what either study could establish alone.[5] Field and survey evidence in Christ et al. (2025), in which

[5]The two samples also weight brokerage types differently—Shanthikumar and Yoo (2026) study a set of 16 large banks, while our panel spans 79 brokerages of varying size—so the estimated benefits of AI adoption need not coincide across the two studies. Our cross-sectional evidence (Section V) indicates why such heterogeneity arises: the benefit of brokerage AI concentrates where forecast accuracy turns on processing public, machine-readable information and among analysts receptive to algorithmic tools.

analysts describe automating routine work and relying on relational networks to differentiate, corroborates this pattern in what analysts report; we provide large-sample behavioral evidence on the same reallocation margin.

Third, we trace the consequences of institutional AI for the market value and organization of analyst research. Research produced by analysts at more AI-intensive brokerages is priced as more informative, extending the disclosure-processing-cost framework of Blankespoor, deHaan, and Marinovic (2020) from firms' disclosure technologies to intermediaries' production technologies. We also show that AI shifts the experience composition of the analyst workforce, linking institutional AI adoption to the production of research and to the organization of sell-side labor.

The remainder of the paper proceeds as follows. Section II reviews the related literature and develops our hypotheses. Section III describes the AI-adoption measure and the sample. Section IV examines the determinants of adoption. Section V presents the effect on forecast accuracy and its cross-sectional variation. Section VI develops the mechanism. Section VII examines the capital- and labor-market consequences. Section VIII concludes.

II. RELATED LITERATURE AND HYPOTHESES DEVELOPMENT

AI and Analysts' Forecast Accuracy

Sell-side analysts are central to financial markets, shaping investor expectations and influencing asset prices through their earnings forecasts, which inform valuation models, investment strategies, and corporate decisions (Brown and Rozeff 1978; Frankel and Lee 1998; Fried and Givoly 1982). Yet, analysts' forecast accuracy is limited by cognitive constraints and time pressures when processing complex information (Driskill, Kirk, and Tucker 2020; Hirshleifer, Lim, and Teoh 2009; Hirshleifer et al. 2019; Jiao 2024). These challenges have intensified due to the rapid growth in information volume and complexity, as analysts navigate vast amounts of corporate disclosures, news, and alternative data sources. The resulting complexity strains analysts' cognitive capacity, potentially leading to delayed or suboptimal forecasts (Hodder, Hopkins, and Wood 2008; Lehavy, Li, and Merkley 2011).

Against this backdrop, AI has emerged as a promising tool to enhance analysts' information-processing capacity. By automating tasks such as data extraction, modeling, and pattern recognition, AI enables more efficient information acquisition and analysis, with the potential to improve forecast accuracy. A growing literature explores AI's role in financial forecasting, often through horserace tests comparing machine-learning models to human analysts. For example, Ball and Ghysels (2018) show that machine-learning models can rival or outperform analysts in predicting firm-level earnings, particularly with high-dimensional data. van Binsbergen et al. (2023) build a machine-learning model of the term structure of earnings expectations that forecasts more accurately than analysts and exposes predictable, conditional biases in their forecasts. Similarly, Cao et al. (2024) find that an AI model integrating textual, financial, and macroeconomic data outperforms analysts, especially for firms with structured disclosures. While these studies do not directly assess AI adoption by analysts, they underscore AI's potential to enhance forecast accuracy.

Parallel research begins to provide initial empirical evidence on the role of AI in shaping financial intermediaries' information production. Coleman et al. (2022) show that AI-generated stock recommendations carry economically meaningful investment value, indicating AI's potential to independently generate useful market insights. Blankespoor, deHaan, and Zhu (2018) find that algorithm-generated earnings articles increase trading volume and liquidity, especially among retail investors. In a more recent study, Bertomeu, Lin, Liu, and Ni (2025) exploit Italy's temporary ban of ChatGPT as a natural experiment and show that the ban raised forecast dispersion among local analysts, suggesting AI may aid analysts' information processing. While these studies highlight the value of AI-generated information, they treat AI as either a standalone source of earnings forecasts or a general-purpose technology, rather than examining the effects of its systematic integration into analysts' workflows or in reshaping the information production process within financial institutions. A related but distinct margin is the adoption of big data. Chi, Hwang, and Zheng (2025) show, from the text of analysts' own reports, that individual analysts have begun drawing on alternative data and that forecasts informed by such data are more accurate. Alternative data, however, are an information input, not a processing technology. They expand what an analyst can process, whereas AI supplies the processing capability. The adoption margin also differs—theirs is the individual analyst's choice of inputs, ours the brokerage's institutional investment in processing capability. If AI adoption improves analysts' ability to generate accurate forecasts, we expect to observe a positive association between brokerage AI adoption and forecast accuracy. Accordingly, we formulate our first hypothesis:

H1: *AI adoption is positively associated with the relative accuracy of analysts' earnings forecasts.*

The foregoing discussion notwithstanding, it is not obvious that the predicted association will hold. The productivity implications of AI adoption have been widely studied in the broader literature on technology and organizational change. A consistent theme is that AI and similar technologies do not simply substitute for human labor but rather interact with organizational processes and human capital to jointly shape outcomes (Aral, Brynjolfsson, and Wu 2012; Lindebaum, Vesa, and den Hond 2020). The outcome of such interactions is inherently uncertain and context-dependent, as the integration of AI into workflows often requires substantial organizational adaptation and creates hybrid systems where human and AI-generated information are jointly processed. Moreover, prior research shows that the productivity gains from information technology depend on its actual use within the organization and on complementary investments in skills, processes, and infrastructure (Aral et al. 2012; Devaraj and Kohli 2003). Taken together, these insights suggest that the effects of AI adoption depend on organizational synergies and the hybridization of human and AI inputs.

Public-Information Dependence, Accessibility, and Analyst Receptivity

A foundational insight from the computer science and information systems literature is that AI models are inherently constrained by the availability and accessibility of their data input. Unlike humans, who can interpolate, contextualize, and draw on tacit knowledge, the performance of AI systems depends heavily on the availability and scale of observable, machine-readable data (Domingos 2012; Halevy, Norvig, and Pereira 2009). As a result, even sophisticated AI models require access to large volumes of well-labeled and semantically clear data to perform effectively (Jordan and Mitchell 2015). In applied domains such as healthcare or legal analytics, this limitation manifests in well-documented ways. While AI systems perform well when data are standardized and codified, they struggle when faced with ambiguity, inconsistency, or informational gaps (Marcus and Davis 2019). These observations highlight the critical importance of having abundant, publicly accessible, and usable data to support robust algorithmic outputs.

Analysts typically draw on both public and private information to make earnings forecasts. Prior work shows that analysts frequently gain value-relevant insights through relationship-based interactions with corporate managers at conferences, earnings calls, and investor meetings (Green, Jame, Markov, and Subasi 2014a,b; Kirk and Markov 2016; Mayew, Sharp, and Venkatachalam 2013). These interactions often yield soft information, such as managerial tone, informal guidance, or off-the-record cues, that are not captured in public disclosures but materially influence analysts' forecasts. Accordingly, a long stream of research finds that short-horizon analyst forecasts often outperform econometric models because of access to such non-public cues (Bradshaw, Drake, Myers, and Myers 2012; Brown and Rozeff 1978; de Silva, Thesmar, and Giglio 2024).

Because AI primarily enhances the processing of public information, its differentiating effect is most likely to emerge under two conditions. First, the forecast setting must favor the processing of public information. Second, the analyst must be willing to rely on the tools that process it. The first condition has two requirements. Public information must play a dominant role in forecasts, so that variation in forecast accuracy depends more on how common data are interpreted than on privileged access to private signals. And the information environment must be conducive to algorithmic extraction, with structured, comprehensible inputs available in formats AI can parse.

We capture the first requirement along two contextual dimensions. First, we consider forecasts issued earlier in the forecasting cycle, when analysts must rely more on historical financials and macroeconomic indicators due to the lack of updated guidance or interactions with management. Second, we examine settings with lower information asymmetry among analysts, where they work from more comparable information sets and public disclosures are more likely to contain the bulk of value-relevant information. In both cases, public information plays a greater role in driving forecast accuracy, which reduces the influence of information access disparities and shifts the source of variation in analysts' forecast accuracy toward differences in processing ability. As a result, AI's processing advantage is more likely to translate into superior forecast accuracy.

To capture the second requirement, we consider the readability of firm disclosures as a proxy for the accessibility and extractability of public information.[6] Even when analysts have access to similar disclosures, variation in linguistic complexity can substantially affect the ease with which value-relevant information is identified and processed, particularly by algorithmic tools. Readable disclosures are more structured, less ambiguous, and thus more amenable to algorithmic parsing, enabling AI tools to more effectively extract patterns, cues, and signals. Thus, disclosure readability may shape the effectiveness of AI in enhancing analysts' forecasting performance. Building on this idea, we expect the positive association between AI adoption and forecast accuracy to be more pronounced in contexts where forecast accuracy depends more heavily on analysts' information-processing capabilities and where information is more accessible to algorithmic systems.

The second condition lies not in the firm's information but in the analyst. Realizing AI's benefit requires that analysts rely on the tools their brokerage provides. However, algorithm aversion—the tendency to discount algorithmic advice (Dietvorst et al. 2015)—can inhibit such reliance and varies across individuals and with the nature of the task (Castelo, Bos, and Lehmann 2019). Analysts with technical backgrounds, being more familiar and comfortable with such tools, are therefore likelier to engage with AI and to extract value from it. In the task-based terms we develop below, the first condition identifies settings in which the forecasting subtask mix tilts toward routine, machine-processable work; the second identifies the analysts most likely to realize AI's value in that work.

H2: *The positive relation between AI adoption and relative forecast accuracy is stronger when forecasts are issued earlier in the forecasting cycle (H2a), when information asymmetry among analysts is lower (H2b), when corporate disclosures are more readable (H2c), and when analysts*

[6] Although readability is traditionally defined from the perspective of human comprehension, prior research shows that texts rated as highly readable by humans tend to align closely with how AI models parse and extract meaning, reflecting a strong correlation between human and machine readability. For instance, Trott and Rivière (2024) report that GPT-4's zero-shot readability ratings correlate highly with human judgments.

have technical backgrounds (H2d).

Mechanism: Within-Task Substitution and Between-Task Complementarity

A natural question is how brokerage AI adoption improves the forecasts of the analysts it employs. We organize our predictions around the task-based view of automation and human-AI collaboration (Acemoglu 2025; Acemoglu and Restrepo 2018, 2019; Autor 2015; Fügener et al. 2026). In this framework, production is organized into tasks that can be performed by human labor, by capital such as software and algorithms, or by humans working with AI. *Automation* occurs when AI performs tasks or subtasks humans would otherwise perform; *augmentation* occurs when AI works with a human on the same task by providing advice, structure, or decision support; and *reallocation* occurs when the human capacity released by automation is redirected toward other tasks.

This framework fits our setting because an earnings forecast is not a single indivisible task but a bundle of subtasks, including information retrieval, disclosure parsing, model updating, peer comparison, anomaly detection, scenario analysis, the assessment of management credibility, and final judgment. Some subtasks are routine, codifiable, and machine-processable; others require contextual interpretation, private information, and relationship-based knowledge. Within the forecasting task, brokerage AI can therefore play one of two roles. It may automate routine subtasks, substituting for the analyst's cognition and conserving her scarce cognitive resources, or augment the analyst on judgment-intensive subtasks, a within-task complementarity whose value depends on the judgment she still supplies (Fügener et al. 2026). These roles are not mutually exclusive, and our first question is which dominates. Forecast production is, however, only one of the analyst's tasks. To the extent AI automates routine subtasks and frees cognitive capacity, the analyst can reallocate that capacity to other tasks in which human input retains a comparative advantage over AI—a between-task complementarity (Acemoglu 2025; Fügener et al. 2026). Our second question is whether such reallocation occurs. This second question presupposes the first. Reallocation arises only if automation, rather than augmentation, dominates within the forecasting task, since only automation frees capacity to redeploy.

Human analysts operate under time pressure, bounded attention, and cognitive constraints (e.g., Driskill et al. 2020; Hirshleifer et al. 2009; Jiao 2024), and recent work shows they are prone to decision fatigue, a phenomenon in which, as an analyst issues more forecasts within a single day, accumulated cognitive depletion degrades the quality of later judgments (Hirshleifer et al. 2019; Jiao 2024). Decision fatigue separates automation from within-task complementarity, because the two imply opposite patterns as the analyst's cognition is depleted. If AI automates routine subtasks, it substitutes for the analyst's cognition, so its marginal value is greatest precisely when that cognition is scarcest, and its benefit should be larger for fatigued forecasts. If instead AI works mainly by augmenting the analyst's judgment within the task, its value depends on the judgment the analyst can still supply, which fatigue erodes, so its benefit should not increase, and may fall, for fatigued forecasts. The two forces need not be mutually exclusive, and nothing in these arguments establishes ex ante which dominates, so we state the hypothesis non-directionally and let the fatigue margin, where the analyst's own cognition is most constrained, adjudicate. A benefit that is larger under fatigue would indicate that the cognition-conserving automation role dominates; a benefit that does not increase, or falls, would indicate that the judgment-dependent augmentation role dominates.

H3a: *The association between AI adoption and relative forecast accuracy differs between fatigued and non-fatigued forecasts.*

Herding toward the consensus is a well-documented analyst heuristic (Clement and Tse 2005; Hong, Kubik, and Solomon 2000; Welch 2000), and decision fatigue can heighten reliance on such low-effort heuristics (Hirshleifer et al. 2019). Herding therefore provides a second, behavioral signature of the same within-task mechanism. Under the automation channel, AI conserves the cognition that fatigue depletes, supporting independent and systematic processing precisely when analysts would otherwise default to the consensus, so fatigue-induced herding should be attenuated by brokerage AI adoption. Under the augmentation channel, any reduction in herding operates through the judgment AI enhances, which is the very input fatigue erodes, so the reduction should not strengthen, and may weaken, for fatigued forecasts. Because the two channels again imply opposite signs at the fatigue margin, we state this hypothesis non-directionally as well: attenuation of fatigue-induced herding would corroborate the automation reading of H3a, whereas its absence would be consistent with the augmentation reading.

H3b: *The association between AI adoption and herding toward the consensus forecast differs between fatigued and non-fatigued forecasts.*

The capacity AI conserves need not only raise the accuracy of a given forecast; it can be reallocated. In the task-based view, a technology that takes over routine tasks raises the value of the tasks that remain and frees workers to perform them (Acemoglu and Restrepo 2018, 2019; Autor 2015). A concrete precedent is the bank teller after the automated teller machine (Autor 2015; Bessen 2015). By automating routine cash handling, ATMs did not eliminate tellers but shifted them toward relationship banking. The analyst's case is similar in form. As AI absorbs routine public-information processing, two forces push analysts toward relationship-based work. On the supply side, automation frees the analyst's cognitive capacity; on the demand side, AI erodes the advantage she earns from processing public information, a task it increasingly performs as well as she does. Both point toward relationship building—cultivating access to firm management, a long-documented source of analysts' informational advantage (Green et al. 2014a; Mayew 2008), and standing with the institutional clients whose assessments shape analyst careers (Yezegel 2023)—work that turns on the analyst's own effort and judgment and that AI cannot perform for her. Earnings calls offer an observable window into this effort. Engagement on a call is a single activity with two audiences. It deepens the analyst's relationship with management, which exercises discretion over which analysts are recognized in the question-and-answer session and in what order (Cohen et al. 2020), and it displays her preparation and standing to the clients listening in. If AI frees capacity that analysts reinvest in these relationships, analysts at more AI-intensive brokerages should engage more actively on these calls.

H3c: *AI adoption is associated with greater analyst investment in relationship building, reflected in higher earnings-call participation, an earlier position in the question-and-answer queue, and greater information elicitation from management.*

III. AI-ADOPTION MEASURE AND SAMPLE CONSTRUCTION

Measuring AI Adoption from Job Postings

A central empirical challenge is that analysts' use of AI is rarely observable at the forecast level. The tools, models, data infrastructure, and research platforms that support analysts' information processing are typically proprietary and embedded within brokerage-level workflows (e.g., Brown, Call, Clement, and Sharp 2015). Following Babina et al. (2024) and a growing literature that measures firm AI activity from online job postings (e.g., Acemoglu, Autor, Hazell, and Restrepo 2022; Goldfarb, Taska, and Teodoridis 2023), we therefore proxy for analysts' AI adoption using brokerage-level demand for AI skills in job postings for analyst research-related positions.[7] The measure reflects how much a brokerage invests in AI capabilities that support analysts' information acquisition, processing, and forecasting.[8]

We obtain job-posting data from Lightcast (formerly Burning Glass Technologies), which provides broad coverage of online postings in the United States beginning in 2010.[9] Each posting records the employer, position title, required skills, and posting date. We retain full-time postings for positions in analysts' research workflow, identified from job titles and descriptions. These positions include frontline analyst roles, such as equity analysts and financial analysts, and adjacent research roles, such as research scientists and quantitative researchers.[10]

For each brokerage-quarter, we define *AI Hiring* as the share of full-time research-related job postings that require at least one AI-related skill, namely artificial intelligence, natural language

[7]Throughout the paper, "analysts' AI adoption" is shorthand for analysts being supported by AI capabilities implemented at the brokerage level.

[8]Proxying for a firm's investment in a technology by its demand for the skills that build and operate it is a standard approach, because such capabilities are embodied largely in specialized human capital rather than in separately observable capital expenditures (e.g., Tambe 2014 for big data; Babina et al. 2024 for AI). Recovering that skill demand from job-posting data, as we do, is likewise well established (Deming and Kahn 2018; Ham, Hann, Rabier, and Wang 2024).

[9]We obtained the Lightcast data through an institutional license purchased in early 2023.

[10]This restriction serves two purposes: it captures within-organization complementarities that arise when analysts are supported by adjacent research personnel, and it excludes AI investments unlikely to affect the analyst function, such as general information-technology roles ("network engineer," "cloud architect," "software developer"). In untabulated tests, both a narrower measure based solely on frontline analyst roles and a broader measure imposing no role restriction (the AI share across all of a brokerage's postings) yield highly consistent results.

processing, or computer vision:

$$AI\ Hiring_{b,t} = \frac{\text{\# of AI-related research postings}_{b,t}}{\text{\# of research-related postings}_{b,t}}, \tag{1}$$

where b indexes brokerages and t indexes calendar quarters. Figure 1 plots the evolution of *AI Hiring* in Panel A.[11]

We then measure brokerage-level AI adoption as the average AI-hiring share over the prior four quarters:

$$AI\ Adoption_{b,t} = \frac{1}{4}\sum_{s=1}^{4} AI\ Hiring_{b,t-s}. \tag{2}$$

The backward-looking window allows for recruitment, onboarding, and integration lags, and it smooths transitory hiring shocks, particularly in the presence of employee turnover.[12] The measure shows substantial cross-brokerage variation in AI-related hiring intensity, reaching more than 31 percent of research-related postings in the single most AI-intensive brokerage-quarter; averaged over a brokerage's sample quarters, even the most AI-intensive brokerage reaches about 4 percent.[13] These adoption differences are also persistent. The average adjacent-year Spearman correlation of brokerages' AI-adoption ranks is 0.83, about two-thirds of top-quartile brokerages remain in the top quartile the following year, and the correlation between a brokerage's average

[11] Panel B plots the share of non-AI technology hiring. The two series differ sharply in both level and dynamics: AI hiring remains a small share of research-related postings, rising mainly before the COVID-19 pandemic and peaking below 2 percent, whereas non-AI technology hiring is an order of magnitude larger and rises steadily from roughly 11 percent to 26 percent over the sample period. This divergence indicates that AI hiring is not a simple relabeling of broad technology hiring within brokerage research departments.

[12] Measuring firm-level labor demand from job-posting flows accumulated over a trailing window is well established (e.g., Deming and Kahn 2018; Ham et al. 2024). The four-quarter horizon is a design choice intended to capture the recent organizational investments most likely to affect analysts' current behavior, and our inferences do not hinge on it: re-estimating the main specification with longer trailing windows of eight, twelve, and sixteen quarters, and with an alternative *stock* measure (the cumulative AI share of a brokerage's hiring from 2010Q1 through quarter $t-1$), leaves the positive and significant association between AI adoption and forecast accuracy unchanged (Table 5, Panels B and C). A cumulative stock is a natural proxy for the *level* of accumulated AI capability, but three considerations favor the trailing flow. First, AI capability in this setting is embodied in specialized human capital and so depreciates (employees turn over and the skill content of AI changes quickly), which a non-depreciating stock ignores; the trailing window is, in effect, a stock with rapid depreciation that weights the recent investments most likely to shape analysts' current research environment. Second, comparable stocks require a complete posting history from 2010Q1 for every brokerage, which reduces the estimation sample by roughly a fifth, whereas the flow measure retains the full sample. Third, a cumulative average becomes increasingly inertial as history accumulates, leaving little identifying variation within analyst-firm pairs, the very variation on which our most demanding specifications rely.

[13] Online Appendix Table OA1 lists the ten brokerages with the highest trailing-window, quarter-averaged AI adoption in our sample.

rank over 2012–2016 and over 2017–2022 is 0.70.

An alternative construction would cumulate past hiring into a stock. We view the trailing flow as the better proxy for the AI capability analysts currently draw on, for three reasons. First, AI capability is embodied in specialized human capital and depreciates, through employee turnover and through the fast-changing skill content of AI itself; a non-depreciating stock ignores this depreciation. Second, stocks are comparable across brokerages only when every brokerage's hiring history is observed over the same window, which requires restricting the sample to brokerages with complete posting histories and costs roughly a fifth of the forecast sample. Third, a cumulative average becomes increasingly inertial as the sample proceeds, leaving little identifying variation within analyst-firm pairs.[14]

This proxy has three advantages. First, it links brokerage-level AI investment to a large sample of individual analyst forecasts in I/B/E/S. Second, it captures AI as an input that supports human analysts rather than as a standalone forecasting technology or as an algorithmic recommendation that analysts may choose to override. Third, by measuring adoption at the brokerage level, it reflects organizational spillovers and shared research infrastructure that can affect the behavior and performance of individual analysts within the same brokerage.

The measure is also deliberately conservative. Brokerages can adopt AI through channels not reflected in hiring data (outsourcing, off-the-shelf services, or public tools). However, brokerages that hire AI talent into research-supporting roles are making among the larger and more durable commitments to AI. The inclusion of brokerages that adopt AI through lighter channels in the comparison group blurs the distinction between more- and less-AI-intensive brokerages, biasing our estimates toward zero and understating the full effect of institutional AI.

Sample Selection

A critical step is to identify each analyst's brokerage at the time of each forecast and match it to the corresponding Lightcast employer. We begin with the 98 brokerages with EPS forecasts on

[14] A stock version of the measure nonetheless delivers the same inference (Table 5, Panel C), and the lag structure of the baseline relationship supports the flow interpretation: the accuracy benefit tracks recent rather than long-past hiring (Online Appendix Table OA9).

I/B/E/S in every year from 2018 through 2022.[15] We recover each forecast's brokerage affiliation from I/B/E/S and match brokerage names to Lightcast employer names using automated fuzzy matching and manual verification.[16] This procedure matches 79 of the 98 brokerages to Lightcast employers (Table 1, Panel A), which form the core sample for our brokerage-level analyses.[17]

Our main tests are conducted at the forecast level. We collect one-year-ahead EPS forecasts issued by analysts at these 79 brokerages from the I/B/E/S Detail History file over 2012 through 2022.[18] Following prior literature, we exclude forecasts that are missing firm identifiers or required controls, as well as forecasts for firms in the utilities and financial industries (SIC 4900–4999 and 6000–6999). Because our main dependent variable measures forecast accuracy relative to peers covering the same firm in the same year, we also drop forecasts in firm-years followed by fewer than two analysts. This selection procedure yields a final sample of 816,418 forecasts (Table 1, Panel B).[19] These forecast-level tests use the main forecast sample, or subsamples of it, combined, where relevant, with additional data from Compustat for accounting variables, CRSP for stock returns, S&P Capital IQ for earnings-call transcripts, the SEC Analytics Suite for firm disclosure readability, Revelio Labs for analyst educational backgrounds, and the Google COVID-19 Community Mobility Reports for county-level workplace mobility.

[15] We require activity from 2018 through 2022, rather than throughout the full 2012–2022 sample period, so as not to exclude brokerages that first appear after 2012; AI-related hiring also became more prevalent after 2018, and a shorter continuity window reduces merge errors across data sources with inconsistent historical coverage. This choice retains 13 additional brokerages, such as Melius Research, that account for only a small share of forecasts: the remaining 66 continuously active brokerages still issue 97.8 percent of the sample (798,281 of 816,418 forecasts). Online Appendix Table OA4 re-estimates the baseline on this continuously active subsample and obtains nearly identical estimates.

[16] Because the I/B/E/S Detail History file does not report brokerage names, we link it to the I/B/E/S Recommendation file using company identifiers, analyst codes, and forecast timing. We then construct a dictionary mapping each broker's Estimator ID, or "ESTIMID," to its analysts' names, the initial of the first name plus the full last name, and the corresponding announcement dates. Using FINRA's BrokerCheck and LinkedIn, we manually identify the brokerage each analyst was affiliated with when issuing the forecast and recover its full name.

[17] The matched brokerages are, on average, larger than those in the full I/B/E/S universe. Because roughly 49 percent of forecasts are issued by analysts at the largest 10 percent of brokerages by headcount, the matched brokerages account for the bulk of forecast activity, although our inferences may be less representative of the smallest brokerages.

[18] The sample begins in 2012 because the Lightcast data begin in 2010 and our construction requires a lagged posting history. The sample ends in 2022, the last year with complete job-posting coverage.

[19] Because each test imposes its own data requirements, the applicable sample is often smaller than the main sample; we therefore introduce only the main forecast sample here and report the exact number of observations and any additional restrictions alongside each analysis.

Other Key Variables

Our main dependent variable is analysts' relative forecast accuracy. Following prior literature (e.g., Clement 1999; Jacob, Lys, and Neale 1999), we benchmark each analyst's absolute forecast error against the distribution of forecast errors among analysts covering the same firm in the same fiscal year:

$$\textit{Relative Accuracy}_{i,j,t} = \frac{\text{Median}\left(FE_{\cdot,j,t}\right) - FE_{i,j,t}}{\text{SD}\left(FE_{\cdot,j,t}\right)}, \tag{3}$$

where i indexes analysts, j indexes firms, t indexes fiscal years, $FE_{i,j,t}$ is the absolute value of the difference between analyst i's EPS forecast for firm j and realized EPS for fiscal year t, and the median and standard deviation are computed across all analysts covering firm j in fiscal year t. Higher values of *Relative Accuracy* indicate more accurate forecasts relative to peers. This within-firm-year benchmarking mitigates the influence of firm characteristics and time trends and isolates differences in analyst performance within a narrow peer set.

Our key independent variable is *AI Adoption*, the brokerage's trailing four-quarter AI-hiring share defined in Equation (2). For each forecast, we assign the value of *AI Adoption* for the analyst's brokerage in the quarter corresponding to the forecast date. The control variables follow prior research on analyst forecast accuracy and include firm-specific experience (*Firm Experience*), general experience (*General Experience*), forecast effort (*Effort*), broker size (*Broker Size*), forecast age (*Forecast Age*), the number of firms the analyst follows (*Firms Followed*), and the number of analysts covering the firm (*Analyst Coverage*). We min-max standardize *AI Adoption* and all controls within each firm-year, so each regressor is measured relative to the distribution of analysts and brokerages covering the same firm in the same year. This scaling places variables on a common unit interval and applies the same within-firm-year grouping used to construct the dependent variable, consistent with the within-firm-year scaling in Hirshleifer et al. (2019). All variables are defined in Appendix A, and Table 2 reports descriptive statistics.

Validating the AI-Adoption Measure

We validate the AI-adoption measure against two independent indicators of brokerage AI activity. First, brokerages with greater AI hiring also produce more AI innovation. Following

Babina et al. (2024), who relate a labor-based AI measure to innovation output, we identify each brokerage's AI-related patents from USPTO filings and relate them to the hiring-based *AI Adoption* measure. Table 3, Panel A reports a positive and statistically significant association between *AI Adoption* and both the *quantity* of AI patents (measured using the count and share of AI-related filings, Columns 1–4) and the *quality* of AI patents (measured using citation-weighted and cohort-adjusted forward citations, Columns 5–8). The association is present both with and without brokerage fixed effects and after controlling for brokerage characteristics.[20]

Second, brokerages' AI hiring is also positively related to their AI-related disclosures. For the publicly listed brokerages in our sample, we construct *AI Disclosure*, which captures the intensity of artificial-intelligence discussion by a brokerage's own management in earnings calls and investor communications.[21] Table 3, Panel B shows that *AI Disclosure* is strongly and significantly related to the hiring-based *AI Adoption* measure, both with and without brokerage fixed effects and after controlling for brokerage characteristics.[22]

Taken together, these validation tests show the hiring-based measure is positively associated with both AI patenting and public AI disclosure, two indicators measured independently of the Lightcast data. The evidence supports interpreting *AI Adoption* as a proxy for economically meaningful brokerage investment in AI capabilities that can support analyst research.

IV. DETERMINANTS OF BROKERAGES' AI ADOPTION

Design

We start by examining which brokerages adopt AI and under what conditions. The brokerage-level setting allows us to address two related questions. First, we characterize the economic forces associated with AI adoption in sell-side research, with particular attention to diffusion across competing brokerages. Second, to distinguish AI adoption from general digitization, we test whether the same forces also explain non-AI technology hiring.

[20] Online Appendix Figure OA1, Panel A shows the same positive relation in a binned scatter plot.

[21] Using firms' narrative disclosures to infer their underlying activities and attributes is well established in the accounting literature (e.g., Li, Lundholm, and Minnis 2013; Merkley 2014).

[22] Online Appendix Figure OA1, Panel B plots the same positive relation.

We frame AI adoption as a cost-benefit decision.[23] Four sets of factors shape this trade-off: the brokerage's competitive environment, its organizational capacity, its analysts' human capital, and the demand for research implied by its coverage portfolio. This organization parallels the grouping of analyst-level determinants into ability, resources, and portfolio complexity in Clement (1999). Peer adoption can raise the perceived benefits of AI by revealing its value and raising the competitive cost of delay (DiMaggio and Powell 1983; Hall and Khan 2003). Brokerage capacity affects the ability to absorb the costs. Larger brokerages have more resources and scale economies, while broader research operations may face greater integration frictions. Analyst human capital eases implementation. Specifically, experienced analysts are better positioned to evaluate new tools. Finally, a brokerage's coverage portfolio proxies for the demand for differentiated research. AI is more valuable where covered firms attract sophisticated investors and sharper analysis pays off more. We estimate the following brokerage-quarter specification:

$$AI\ Hiring_{b,t} = \alpha + \beta Peer\ Adoption_{b,t} + \gamma BC_{b,t} + \delta\, HC_{b,t} + \lambda\, PD_{b,t} + \phi + \varepsilon_{b,t}, \tag{4}$$

where b indexes brokerages, t indexes calendar quarters, and ϕ denotes the fixed effects included in the corresponding specification (none, year, or year-quarter). The dependent variable, *AI Hiring*, is defined in Equation (1). The competitive environment is captured by *Peer Adoption*, the number of other brokerages in b's size quantile with at least one AI-related hire before quarter t.

Brokerage-level capacity, $BC_{b,t}$, includes broker size (*Broker Size*), the number of firms covered (*Firm Coverage*), and the number of unique industries covered (*Industry Coverage*). Analyst human capital, $HC_{b,t}$, is proxied by the brokerage's average analyst experience (*Mean Experience*). Portfolio-level demand for analyst research, $PD_{b,t}$, is measured using characteristics of the brokerage's covered firms: average institutional ownership (*Inst Ownership*), the number of unique institutional holders (*# of Institutions*), average analyst coverage (*Mean Coverage*), forecast dispersion (*Forecast Dispersion*), intangible intensity (*Intangibles*), return volatility (*Std Ret*), market capitalization (*MVE*), and market-to-book ratio

[23] The expected benefits include improvements in research production, analyst productivity, client retention, and competitive positioning. The costs include acquiring specialized talent, integrating new tools into existing research processes, and adapting organizational routines.

(*Market to Book*).[24] Variable definitions appear in Appendix A.

To assess whether these determinants are specific to AI, rather than common to technology investment more broadly, we estimate a parallel specification using the brokerage's non-AI technology hiring, *Tech Hiring*, as the dependent variable.[25]

Results

Table 4 presents the results.[26] First, AI hiring is strongly related to the competitive environment. The coefficient on *Peer Adoption* is positive and highly significant across specifications, indicating that brokerages hire more AI-related talent when more of their size-group peers have already done so. This pattern is consistent with technology diffusion (e.g., Griliches 1957; Mansfield 1961; Rogers 2003) as peer adoption may reveal AI's value, increase the perceived cost of delay, or both.

Second, organizational capacity helps explain adoption. Larger brokerages and brokerages covering more firms invest more in AI, consistent with resource availability and scale economies in technology adoption. By contrast, *Industry Coverage* is negatively associated with AI hiring, suggesting broader research operations face greater integration and coordination costs. Analyst human capital also matters. AI hiring is higher when the brokerage's analysts have more experience on average, consistent with experienced teams better able to absorb and deploy new technology.

Third, AI hiring varies with the expected demand for differentiated research. Brokerages covering firms with higher institutional ownership and more institutional holders invest more in AI, consistent with larger returns to sharper analysis when coverage portfolios are institutionally important. AI hiring is lower among brokerages whose portfolios tilt toward larger and higher market-to-book firms, suggesting weaker incremental demand for AI-supported analysis in settings

[24] For the determinants analysis, we standardize all variables in Equation (4) using Z-scores prior to estimation and winsorize all continuous variables at the 1 percent level.

[25] We use the same regressors as in Equation (4), with the sole exception of *Peer Adoption*: non-AI technology hiring is already widespread throughout the sample, so a peer-diffusion regressor is neither well defined nor informative.

[26] Across the main test and our robustness tests, we report results under a range of specifications, with and without fixed effects, to triangulate the documented associations, following Armstrong, Kepler, Samuels, and Taylor (2022). In our remaining tests, we adopt the most demanding specification, with analyst-firm fixed effects.

where the marginal payoff to additional research differentiation is lower.[27]

Taken together, AI hiring varies systematically with peer adoption, organizational capacity, analyst human capital, and the expected value of research differentiation.

AI Hiring versus Non-AI Technology Hiring

Table 4, Columns 4–6 repeat the analysis using non-AI technology hiring as the dependent variable. Some scale-related variables predict both outcomes, as expected. The broader pattern, however, differs sharply. The demand-side variables that help explain AI hiring do not carry over to non-AI technology hiring. Institutional ownership, the number of institutional holders, and forecast dispersion are insignificant or oppositely signed, while *MVE* and *Industry Coverage* reverse sign.

This contrast suggests that AI hiring is not simply a relabeling of general technology investment. Non-AI technology hiring appears to track brokerage scale and operational capacity more closely. AI hiring, by contrast, is tied to competitive diffusion and to the expected returns from producing more differentiated analyst research. The distinction reinforces the divergence between the AI and non-AI technology hiring series in Figure 1 and supports the interpretation of AI adoption as a distinct organizational choice.

V. BROKERAGES' AI ADOPTION AND ANALYST EARNINGS FORECASTS

Analyst Forecast Accuracy

To test whether analysts forecast more accurately when their brokerage invests more in AI, we estimate

$$Relative\ Accuracy_{i,j,t} = \alpha + \beta\, AI\ Adoption_{b,t} + \delta\, Controls_{i,j,t} + \phi + \varepsilon_{i,j,t}, \tag{5}$$

where *Relative Accuracy* is analyst i's within-firm-year benchmarked forecast accuracy for firm j in fiscal year t; b denotes analyst i's brokerage at the forecast date, so that $AI\ Adoption_{b,t}$ is the brokerage's trailing four-quarter AI-hiring share; $Controls_{i,j,t}$ collects the analyst- and forecast-level controls defined in Section III; and ϕ denotes the fixed effects included in the corresponding specification (none, analyst, or analyst-firm). Because all regressors are

[27] We interpret these demand-side coefficients descriptively: they indicate where AI investment is more likely to arise, without isolating a single underlying mechanism.

min-max standardized within firm-year, the specification differences out firm-year level effects by construction, and β captures the accuracy gain from being relatively more AI-supported among the analysts covering the same firm in the same fiscal year.

We estimate Equation (5) across all fixed-effects specifications, each estimated without and then with the controls: no fixed effects, analyst fixed effects, and analyst-firm fixed effects. The progression moves from comparisons across analysts, to within-analyst comparisons over time, to within-analyst-firm comparisons that hold fixed the analyst's coverage of a given firm. Standard errors are clustered by analyst throughout.

AI Adoption and Forecast Accuracy: Baseline Results

Table 5, Panel A reports the estimates.[28] *AI Adoption* is positive and statistically significant in every specification. In the most demanding one, with analyst-firm fixed effects and analyst-clustered standard errors (Column 6), the coefficient is 0.025 ($p < 0.01$), implying that, when comparing forecasts by the same analyst on the same firm over time, moving *AI Adoption* from the lowest to the highest value among the analysts covering that firm-year is associated with a 2.5 percent increase in relative forecast accuracy. Given the within-firm-year benchmarking, this is an economically meaningful gain within a tightly defined peer set.[29]

The result holds whether we compare across analysts (Columns 1–2), within analysts over time (Columns 3–4), or within analyst-firm pairs (Columns 5–6). These results support H1: analysts at brokerages that adopt AI more intensively produce more accurate forecasts, consistent with brokerage-level AI investment enhancing analysts' information production.[30]

Alternative Measures of AI Adoption

The association is not specific to how adoption is measured. Panel B of Table 5 re-estimates the baseline with *AI Adoption* measured over trailing windows of eight, twelve, and sixteen quarters, and the positive and significant association holds across all three horizons under both

[28]Online Appendix Figure OA2 plots the within-analyst-firm relationship as a binned scatter.

[29]For comparison, a full-range increase in firm-specific experience raises relative accuracy by 4.6 percent in the same specification, so the association with *AI Adoption* is sizable relative to a leading analyst-level determinant.

[30]The baseline relation is also robust to how forecast accuracy is measured: re-estimating it with the absolute forecast error as the dependent variable under firm-year fixed effects delivers the same inference (Online Appendix Section OA II).

analyst and analyst-firm fixed effects.

Panel C replaces the flow with a *stock* measure that accumulates the brokerage's AI-hiring share from 2010Q1 through the quarter before the forecast, computed for brokerages with complete posting histories so that each enters with the same length of history; the stock measure is likewise positive and significant across all fixed-effects structures, with a coefficient of 0.032 ($p < 0.01$) in the most demanding specification (Column 6). Section III explains why the trailing flow nonetheless remains the primary measure.

Workplace Lockdowns as a Shock to In-House AI Access

We next sharpen the interpretation of the baseline relation by exploiting pandemic workplace lockdowns as a shock to analysts' access to the brokerage's in-house AI infrastructure. AI capabilities in sell-side research are deployed at the brokerage level and embedded in internal research systems, so their value depends on analysts' ability to access that infrastructure; lockdowns restricted access sharply while accumulated AI investment moved slowly, variation that county-level workplace-mobility data are well suited to capture (McLaren and Wang 2020). Because counties entered and exited low-workplace-activity regimes at different points over 2020–2022, we can compare periods in which the same brokerage's AI infrastructure was more versus less accessible. If the accuracy gains operate through analysts' use of in-house AI, they should concentrate when offices are open and attenuate under lockdowns.

We restrict the sample to 2020–2022 and set *Lockdown* to one if workplace visits in the brokerage's headquarters county fell below 60 percent of the January 2020 baseline in the Google COVID-19 Community Mobility Reports during the forecast quarter (Google 2022). This measure varies at the county-quarter level. We re-estimate Equation (5) separately for open and locked-down observations, test the difference in the *AI Adoption* coefficient across the two, and repeat the exercise for each trailing measurement window.[31]

Table 6 reports the results. In open-office quarters, the coefficient on *AI Adoption* is positive

[31] We report the test across progressively longer trailing windows, from four to sixteen quarters, because lockdowns may affect short-run AI hiring. Longer windows ensure that *AI Adoption* primarily reflects accumulated, largely pre-lockdown AI investment rather than hiring decisions made during the lockdown, so that the open-versus-locked comparison isolates variation in access to existing AI infrastructure from lockdown-induced changes in investment.

and highly significant; in locked-down quarters, it falls close to zero and becomes statistically insignificant, and the difference is significant for every trailing window. For the baseline four-quarter window, the coefficient is 0.231 when offices are open and 0.047 when they are locked down ($\chi^2 = 13.07$, $p < 0.001$).

The accuracy gains are thus concentrated when analysts can access the brokerage's in-house AI and largely disappear when lockdowns restrict that access. More than a robustness check, this design moves the evidence closer to a causal interpretation. Because the attenuation arises within the pandemic period and turns on staggered, county-level changes in workplace activity, it is difficult to attribute to a fixed brokerage characteristic correlated with AI adoption or to any single aggregate shock. The evidence is consistent with AI improving forecast accuracy through analysts' use of brokerage-provided AI infrastructure.

Cross-Sectional Variation in the Effect of AI Adoption

We conduct a series of cross-sectional tests to examine the conditions under which brokerage AI is most effective. The hypothesis development in Section II implies two such conditions. First, because AI's contribution operates through the processing of public information, its benefit should be largest where forecast accuracy turns on processing common, machine-readable data rather than on judgment or access to private signals. These include forecasts issued early in the cycle as well as forecasts for firms with lower information asymmetry among analysts or with more readable disclosures (H2a–H2c). Second, because the brokerage supplies the capability but the analyst must choose to rely on it, the benefit should concentrate among analysts with technical backgrounds, who are less prone to algorithm aversion (H2d). The first condition locates where AI's potential is largest while the second, where that potential is realized. We examine each in turn.

The Firm's Information Environment

Brokerage AI should help analysts most when two requirements are met: forecast accuracy must depend on *processing* public information rather than on access to private signals, and that information must be accessible enough for algorithmic extraction. We capture these with three features of the information environment and re-estimate Equation (5) within sample partitions on each (Table 7, Columns 1–6). Because AI's contribution operates through the processing of public data, it should add most where such processing is the binding constraint on accuracy. We expect this to be the case for forecasts made earlier in the period as well as for forecasts of firms where analyst information asymmetry is low or firms where data are structured well enough for algorithms to parse (i.e., where disclosures are more readable). *Forecast Timing* separates forecasts issued in the first half of the forecast cycle, which lean on public and historical information, from later forecasts that can incorporate fresh managerial guidance. *Info Asymmetry* is low when analysts hold comparable information, so that performance turns on processing shared data rather than on private access.[32] *Disclosure Readability*, the reverse-coded Gunning Fog Index of the firm's filings, is higher when disclosures are more amenable to algorithmic parsing.[33]

The effect of *AI Adoption* is concentrated in the predicted subsample in each case: among earlier-cycle forecasts ($\chi^2 = 33.78$, $p < 0.001$), in low-asymmetry settings ($\chi^2 = 18.05$, $p < 0.001$), and for more-readable disclosures ($\chi^2 = 5.64$, $p = 0.018$). These results support H2a–H2c: the benefit of brokerage AI rises with the accessibility of the public information analysts must process, indicating that its value is greater in a more transparent, well-structured information environment.

Analyst Technical Background

The value of brokerage AI may also depend on the analyst's own facility with algorithmic tools. People often discount or avoid algorithmic advice even when it outperforms them, a phenomenon termed "algorithm aversion" (Dietvorst et al. 2015), the strength of which depends on the task (Castelo et al. 2019). We conjecture that it also depends on the user—analysts without a technical background may be both less able to incorporate AI into their workflow and more averse to relying on it, either of which would weaken the benefit of their brokerage's AI. We split

[32] We infer *Info Asymmetry* from the dispersion structure of analyst forecasts following Barron, Kim, Lim, and Stevens (1998), decomposing total forecast-error variance into public and private components and defining asymmetry as the share of uncertainty not attributable to shared information; this isolates private-information differences among analysts more cleanly than raw forecast dispersion.

[33] The readability results are robust to the SMOG and Automated Readability indices.

analysts on technical background, using Revelio Labs profiles to identify those with a STEM education (Table 7, Columns 7–8).[34]

The accuracy gain from brokerage AI is more concentrated among analysts with technical training. Specifically, the coefficient on *AI Adoption* is positive and significant for this group but is insignificant for others, with the difference marginally significant ($\chi^2 = 2.73$, $p = 0.098$). These results are consistent with H2d: because technical background is little associated with accuracy on its own, the pattern points to a complementarity between AI and the skill to use it rather than to technical analysts being better forecasters, with the benefit concentrating among analysts less prone to algorithm aversion.

VI. MECHANISM: EVIDENCE FROM ANALYST DECISION FATIGUE AND RELATIONSHIP BUILDING

To understand how brokerage AI helps analysts, whose work spans forecasting and the tasks around it, Section II poses two questions. The first is within the forecasting task—does brokerage AI make forecasting easier, whether by automating routine subtasks (a substitution benefit that conserves the analyst's cognition) or by augmenting the analyst's judgment (within-task complementarity)? The second is between tasks—if forecasting becomes easier, do analysts reallocate the freed cognition toward other tasks (between-task complementarity)? Figure 2 maps these channels and the decision-fatigue logic that distinguishes them. We answer the first with the decision-fatigue test below, and the second with analysts' relationship building on earnings calls.

Within the Forecasting Task: AI Adoption and Forecasting under Fatigue

We use the decision-fatigue setting, in which forecast quality deteriorates over an analyst's successive same-day forecasts (Hirshleifer et al. 2019; Jiao 2024), to ask which force dominates within forecast production when the analyst's cognition is most strained. We define *Fcast Under Fatigue* as an indicator equal to one if the forecast is the fourth or later forecast

[34] Technical background is an analyst-level attribute and could instead be entered as an interaction with *AI Adoption*; we report a sample split here to keep the presentation consistent with the firm-level partitions above, and an interaction specification produces the same qualitative pattern.

issued by the analyst on the same trading day, and zero otherwise.[35]

This within-analyst-day setting is useful for two reasons. First, fatigue arises only on days when an analyst's cognition is repeatedly taxed by and largely devoted to forecasting, so examining whether brokerage AI helps more as fatigue accumulates isolates its benefit within the forecasting task.[36] Second, comparing an analyst's earlier and later same-day forecasts holds fixed her stable characteristics, brokerage affiliation, and brokerage resources, and the within-firm-year benchmarking of *Relative Accuracy* holds fixed the firm-level information environment; the remaining caveat is that later forecasts may differ in task difficulty, urgency, or sequencing.

As developed in Section II, automation and within-task complementarity imply opposite signs at the fatigue margin (Figure 2, Panel B). Automation predicts a positive interaction between *AI Adoption* and fatigue (wherein AI substitutes for analysts' cognition), and within-task complementarity a nonpositive one (wherein AI complements analysts' cognition). The interaction coefficient captures the net effect across the subtask bundle. A positive (negative) estimate indicates that the cognition-conserving automation (judgment-dependent augmentation) role of AI dominates when analysts are most depleted, not that augmentation (substitution) is absent.

We test these predictions in two steps. First, we re-estimate Equation (5) on the fatigue sample, adding *Fcast Under Fatigue* and its interaction with *AI Adoption*:

$$\begin{aligned} \textit{Relative Accuracy}_{i,j,t} = \alpha + \beta_1 \textit{AI Adoption}_{b,t} + \beta_2 \textit{Fcast Under Fatigue}_{i,t} \\ + \beta_3 \textit{AI Adoption}_{b,t} \times \textit{Fcast Under Fatigue}_{i,t} + \delta\, \textit{Controls}_{i,j,t} + \phi + \varepsilon_{i,j,t}, \quad (6) \end{aligned}$$

where i indexes analysts, j firms, t forecast dates, and b denotes the analyst's brokerage, $\textit{Controls}_{i,j,t}$ collects the baseline controls, and ϕ the fixed effects included in the corresponding specification. Fatigue should reduce accuracy, implying a negative coefficient on *Fcast Under Fatigue*.

[35]Following Jiao (2024), we retain only forecasts time-stamped between 9:00 and 23:00 so that the sequence reflects active intraday forecasting rather than batched overnight entries. Our inferences are unchanged when we define fatigue using the third or fifth forecast of the day instead.

[36]We interpret this test as identifying which of automation and within-task complementarity dominates AI's within-forecast value, not as ruling either out.

If automation conserves cognition, brokerage AI should attenuate the fatigue penalty, implying $\beta_3 > 0$. If AI works mainly through within-task complementarity, $\beta_3 \leq 0$.

Second, we examine herding, a low-effort heuristic on which fatigued analysts may rely. We define *Herding* as one if the analyst's forecast lies between her own prior forecast and the prevailing consensus, and estimate the following conditional logit model:

$$\begin{aligned}\Pr(Herding_{i,j,t} = 1) = f\Big(&\alpha + \beta_1 AI\ Adoption_{b,t} + \beta_2 Fcast\ Under\ Fatigue_{i,t} \\ &+ \beta_3 AI\ Adoption_{b,t} \times Fcast\ Under\ Fatigue_{i,t} + \delta\, Controls_{i,j,t}\Big), \quad (7)\end{aligned}$$

where f is the logistic function and the remaining terms are as in Equation (6). If fatigue increases reliance on consensus-based heuristics, $\beta_2 > 0$. If brokerage AI conserves the cognition fatigue depletes, it should attenuate fatigue-induced herding, implying $\beta_3 < 0$; if it works mainly through within-task complementarity, $\beta_3 \geq 0$.

Table 8, Panel A, reports the results of the accuracy test, which indicate that fatigue is costly. At the lowest *AI Adoption*, a fatigued forecast is 0.037 relative-accuracy units below the analyst's earlier forecasts that day (analyst-firm fixed effects). The interaction is positive and significant across specifications; in Column 3 the coefficient on *AI Adoption* × *Fcast Under Fatigue* is 0.022, offsetting roughly three-fifths of the penalty as *AI Adoption* moves from its lowest to highest value, while *AI Adoption* alone is small and insignificant. The benefit thus concentrates at the fatigued end of the sequence, where cognition is scarcest, rather than lifting accuracy evenly as a uniform productivity gain would. These results support H3a, and the positive sign of the interaction indicates that the cognition-conserving automation role dominates at the fatigue margin.

Panel B reports the herding model. The results indicate that fatigue raises herding. Specifically, in Column 4 the coefficient on *Fcast Under Fatigue* is 0.263, about a 30 percent increase in the odds of herding. The interaction with *AI Adoption* is negative and significant, -0.178, offsetting about two-thirds of that increase as *AI Adoption* moves from lowest to highest. The effect concentrates at the fatigue margin—the unconditional association between *AI Adoption* and herding is insignificant once analyst-firm fixed effects and controls are included. These results support H3b, and the negative sign of the interaction indicates that brokerage AI adoption curbs fatigue-induced herding, the attenuation that only the automation channel predicts.

Together, the two panels point to the same channel: brokerage AI substitutes for part of the cognitive operations analysts would otherwise perform themselves when making forecasts, leaving accuracy less eroded by fatigue and forecasts more independent under it. The next question is whether the cognitive resources conserved within the forecasting task are reallocated toward other, less automatable tasks, above all the relationship building that analysts conduct on earnings calls.

Reallocation of Freed Capacity: Evidence from Relationship Building on Earnings Calls

The fatigue evidence speaks to how AI enters forecast production. We next ask whether the capacity automation conserves is reallocated toward relationship building (between-task complementarity). Cultivating relationships is central to the analyst's role. Access to firm management is a long-documented source of analysts' informational advantage (Green et al. 2014a; Mayew 2008), and standing with the institutional clients her research serves carries substantial weight for analyst careers (Yezegel 2023). We observe this relationship-building effort on earnings conference calls. Engagement on a call is a single activity with two audiences. It deepens the analyst's relationship with management, which exercises discretion over which analysts are recognized in the Q&A and in what order (Mayew 2008; Cohen et al. 2020), and it demonstrates her preparation and standing to the clients listening in. If AI frees capacity that analysts reinvest in this work, analysts at more AI-intensive brokerages should participate more, be recognized earlier, and draw more out of management.[37]

We estimate, at the analyst-firm-quarter level,

$$Y_{i,j,q} = \alpha + \beta\, AI\ Adoption_{b,q} + \delta\, Controls_{i,j,q} + \mu_{j,q} + \eta_{i,j} + \varepsilon_{i,j,q}, \tag{8}$$

where $Y_{i,j,q}$ is one of three outcomes for analyst i on firm j in quarter q, b denotes analyst i's brokerage in quarter q, $Controls_{i,j,q}$ collects the baseline controls, $\mu_{j,q}$ are firm-quarter fixed

[37]Concurrent and independent work by Shanthikumar and Yoo (2026) also examines AI and analysts' information-seeking effort, including participation in earnings calls. We first use within-day decision fatigue to provide evidence that brokerage AI conserves analysts' cognition in routine forecast production; we then ask whether analysts at more AI-intensive brokerages redirect that capacity toward relationship building.

effects, and $\eta_{i,j}$ are analyst-firm fixed effects. Standard errors are clustered by analyst. The analyst-firm fixed effects compare the same analyst-firm pair as brokerage AI adoption changes, and the firm-quarter fixed effects hold fixed the firm-period setting shared by all covering analysts.

The outcomes are whether the analyst participates in the call (*Participation*); conditional on participating, how early and prominently she is called on in the Q&A (*Rel Position* and *Question Position*, for which lower values denote earlier placement); and the standardized length of management's response to her questions (*SRL*; Yezegel 2023), which proxies for how much she draws out of management. Each outcome plausibly reflects both the management-facing and the client-facing dimensions of relationship building, in weights that likely differ across measures; our tests capture the redeployment of analyst effort toward this relationship-based work and do not apportion it between audiences, a decomposition we leave to future research. What matters for our purposes is what the outcomes share. Each turns on the analyst's own preparation, judgment, and real-time engagement, work AI can assist but cannot perform for her.[38]

Table 9 reports the estimates. Analysts at more AI-intensive brokerages are more likely to participate in earnings calls, are called on earlier in the Q&A queue, and elicit longer responses from management. Because these comparisons hold the analyst-firm pair and the firm-quarter fixed, they are difficult to attribute to fixed differences in analyst standing or to firm-period conditions. These results support H3c, and the pattern across all three measures is consistent with analysts redeploying the capacity AI conserves in forecasting toward relationship building.[39]

Taken together with the decision-fatigue evidence, these results describe how the adoption of AI by brokerages reshapes analysts' work. The fatigue tests indicate that AI substitutes for analysts' scarce cognition in routine forecast production and that this automation role dominates augmentation at the fatigue margin within the forecasting task. The earnings-call tests suggest that the capacity conserved through this automation is reallocated toward the relationship building in

[38]To the extent brokerage AI helps the analyst prepare more targeted questions, that assistance operates on a task distinct from routine forecast production, and the analyst still supplies the judgment, relationship capital, and real-time engagement needed to draw out a substantive response.

[39]These improvements in engagement could feed back into forecast accuracy over longer horizons. This poses no conflict with the decision-fatigue results, which are identified from an analyst's within-day forecast rank and concentrate on busy forecasting days, and so are insulated from these slower, between-task channels.

which analysts retain a comparative advantage. Read together, the two tests reveal how brokerage AI improves the quality of analyst research. It automates some of the subtasks within forecasting and helps analysts invest more in the relationship-based tasks that remain difficult to automate.

VII. CAPITAL- AND LABOR-MARKET CONSEQUENCES OF BROKERAGES' AI ADOPTION

Having documented that AI adoption is associated with more accurate forecasts, we examine two consequences: how the capital market values the research of AI-powered analysts, and how brokerages' demand for analyst labor responds.

Capital Markets: Analyst Research Informativeness

If brokerage AI improves the quality of analyst research, the market should react more strongly to forecasts issued by AI-powered analysts. We measure this reaction with the analyst informativeness measure of Frankel, Kothari, and Weber (2006), which scales the average absolute abnormal return on an analyst's forecast days by the firm's total daily absolute abnormal return over the year:

$$Informativeness_{i,j,t} = \frac{1}{N_{i,j,t}} \cdot \frac{\sum_{d \in D_{i,j,t}} \left| AR_{j,d} \right|}{\sum_{d \in T_{j,t}} \left| AR_{j,d} \right|}, \tag{9}$$

where $AR_{j,d}$ is firm j's abnormal return on day d, $D_{i,j,t}$ is the set of $N_{i,j,t}$ days on which analyst i issues a forecast for firm j in year t, and $T_{j,t}$ is the set of firm j's trading days in year t. Higher values indicate that more of the firm's annual price movement is concentrated on the analyst's forecast days. We standardize *Informativeness* within firm-year and estimate the baseline specification with analyst-firm fixed effects.[40]

Table 10, Panel A, reports the results. In Columns 1–2, *AI Adoption* is positively and significantly associated with *Informativeness*, indicating that forecasts by analysts at more AI-intensive brokerages move prices more than those of peers covering the same firm-year. Analysts supported by more brokerage AI thus produce research that is both more accurate and more informative.

[40] Abnormal returns are market-adjusted; size- and industry-adjusted models give similar results.

Columns 3 and 4 test whether this informativeness benefit operates through the accessibility of a brokerage's AI investment to investors. Because AI-related hiring is not itself disclosed, investors can price it only to the extent that they can observe it, and brokerages make it observable through their own communications. Consistent with this logic, Table 3, Panel B shows that brokerages that hire more AI talent also devote a larger share of their earnings-call discussion to AI (*AI Disclosure*). We therefore add *AI Disclosure* and its interaction with *AI Adoption*, estimated on the subsample of publicly listed brokerages for which disclosure is observed. The interaction *AI Adoption* $\times$ *AI Disclosure* is positive and significant, so the informativeness benefit of AI adoption concentrates among brokerages that disclose their AI more. Thus, the market rewards AI-powered analysts more when the brokerage's AI investment is more accessible to investors.

Labor Markets: The Composition of Analyst Hiring

If brokerage AI substitutes for the routine, codifiable components of analysts' work, brokerages that adopt it may change the kind of analysts they hire, even if they do not hire fewer. We test this with brokerage-level hiring data. For each brokerage and year, we classify its analyst-research job postings by experience, using the years of experience required and the seniority of the title, and compute the share that is entry-level, mid-level, and senior.[41] Following the determinants analysis (Section IV), we standardize all variables within year and include the same brokerage- and portfolio-level controls, and we relate each experience share in year $t+1$ to the brokerage's *AI Adoption* in year t:

$$Hiring\ Share^{g}_{b,t+1} = \alpha + \beta\, AI\ Adoption_{b,t} + \delta\, Controls_{b,t} + \varepsilon_{b,t+1}, \tag{10}$$

where b indexes brokerages and t years, and $g \in \{\text{entry}, \text{mid}, \text{senior}\}$ denotes the experience group. $Hiring\ Share^{g}_{b,t+1}$ is the share of brokerage b's analyst postings in group g in year $t+1$, $AI\ Adoption_{b,t}$ is the brokerage's AI-hiring share in year t, and $Controls_{b,t}$ collects the determinant controls from Section IV. All variables are standardized within year, and we estimate the equation separately for each experience group by pooled OLS, with standard errors clustered by brokerage.

[41] A posting is entry-level if it requires two or fewer years of experience, is an internship, or carries a junior title; senior if it carries a senior, vice-president, or director title; and mid-level otherwise.

Table 10, Panel B, reports estimates of Equation (10). A higher AI-hiring share predicts a significantly smaller entry-level share and a significantly larger senior share of next-year analyst postings. This tilt is not a contraction—in levels, AI-adopting brokerages post *more* analyst openings at every experience level (untabulated). Brokerages that invest in AI thus shift the composition of an expanding hiring program toward more-experienced analysts, consistent with AI taking over the routine tasks that junior analysts have traditionally performed while raising the value of the judgment that senior analysts supply. Because the specification, like the determinants analysis, identifies the association from differences across brokerages rather than within a brokerage over time, we interpret it as a structural difference between more- and less-AI-intensive brokerages rather than a sharp within-brokerage shift.

VIII. CONCLUSION

This paper examines how brokerages' adoption of AI, measured from research-related AI hiring, shapes the work of the sell-side analysts they employ. Adoption has risen over the past decade, with cross-sectional variation explained by peer adoption, organizational capacity, analyst human capital, and institutional demand, and it is associated with higher relative forecast accuracy, robustly so across alternative adoption measures and fixed-effects specifications. The gains are more pronounced for forecasts issued earlier in the cycle, for firms with lower information asymmetry among analysts and more readable disclosures, and for analysts with technical backgrounds. Within the forecasting task, brokerage AI substitutes for analysts' scarce cognition in routine forecast production: its benefit concentrates when analysts are most fatigued and is accompanied by less fatigue-induced herding. The capacity this substitution frees appears to be reallocated toward relationship building—cultivating access to firm management and standing with investing clients—where analysts retain a comparative advantage. In capital markets, forecasts by AI-powered analysts are more informative, an effect concentrated among brokerages that disclose their AI; in labor markets, adopting brokerages tilt the experience composition of analyst hiring toward senior roles without reducing total hiring.

Our contributions are threefold. First, we provide a validated, research-specific measure of brokerage AI adoption and a structured account of its determinants. Second, we characterize how human-AI collaboration works in expert forecasting: rather than running analysts against machines, we study analysts working with institutional AI and show that AI substitutes for analysts' cognition in routine forecast production while the freed capacity is redeployed toward relationship building. Third, we trace the consequences of institutional AI for the market value and organization of analyst research, providing, to our knowledge, the first evidence linking brokerage AI adoption to the experience composition of analyst hiring.

Our evidence is associational, and we do not claim to identify the causal effect of AI adoption. Brokerages choose whether and when to invest in AI, and unobserved time-varying factors could co-move with a brokerage's AI hiring and its analysts' performance; designs based on labor mobility are not feasible because our data record job postings rather than individual employees, and too few analysts move between sample brokerages during 2012–2022 to support estimation. Two features of the evidence nonetheless narrow the alternatives. The workplace-lockdown analysis ties the accuracy gains to periods when analysts can access the brokerage's in-house AI infrastructure, using staggered county-quarter variation that is difficult to attribute to fixed brokerage characteristics or a single aggregate shock. The decision-fatigue tests are identified from within-analyst-day variation, holding fixed the analyst, her brokerage, and the firm-level information environment. We therefore read our estimates as robust, mechanism-consistent associations rather than causal effects.

REFERENCES

Acemoglu, D. 01 2025. The simple macroeconomics of ai. *Economic Policy* 40 (121): 13–58. https://doi.org/10.1093/epolic/eiae042.

Acemoglu, D. and P. Restrepo. June 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108 (6): 1488–1542. https://doi.org/10.1257/aer.20160696. https://www.aeaweb.org/articles?id=10.1257/aer.20160696.

Acemoglu, D. and P. Restrepo. 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives* 33 (2): 3–30. https://doi.org/10.1257/jep.33.2.3.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo. 2022. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics* 40 (S1): S293–S340. https://doi.org/10.1086/718327.

Aral, S., E. Brynjolfsson, and L. Wu. 2012. Three-way complementarities: Performance pay, human resource analytics, and information technology. *Management Science* 58 (5): 913–931. https://doi.org/10.1287/mnsc.1110.1460.

Armstrong, C., J. D. Kepler, D. Samuels, and D. Taylor. 2022. Causality redux: The evolution of empirical methods in accounting research and the growth of quasi-experiments. *Journal of Accounting and Economics* 74 (2–3): 101521. https://doi.org/10.1016/j.jacceco.2022.101521.

Autor, D. H. 2015. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives* 29 (3): 3–30. https://doi.org/10.1257/jep.29.3.3.

Babina, T., A. Fedyk, A. He, and J. Hodson. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151: 103745. https://doi.org/10.1016/j.jfineco.2023.103745.

Ball, R. T. and E. Ghysels. 2018. Automated earnings forecasts: Beat analysts or combine and conquer? *Management Science* 64 (10): 4936–4952. https://doi.org/10.1287/mnsc.2017.2864.

Barron, O. E., O. Kim, S. C. Lim, and D. E. Stevens. 1998. Using analysts' forecasts to measure properties of analysts' information environment. *The Accounting Review* 73 (4): 421–433. https://doi.org/10.2308/tar-1325287.

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. https://doi.org/10.1016/j.jacceco.2025.101782.

Bessen, J. 2015. *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth*. New Haven, CT, Yale University Press.

Blankespoor, E., E. deHaan, and C. Zhu. 2018. Capital market effects of media synthesis and dissemination: Evidence from robo-journalism. *Review of Accounting Studies* 23 (1): 1–36. https://doi.org/10.1007/s11142-017-9422-2.

Blankespoor, E., E. deHaan, and I. Marinovic. 2020. Disclosure processing costs, investors' information choice, and equity market outcomes: A review. *Journal of Accounting and Economics* 70 (2–3): 101344. https://doi.org/10.1016/j.jacceco.2020.101344.

Bradshaw, M. T., M. S. Drake, J. N. Myers, and L. A. Myers. 2012. A re-examination of analysts' superiority over time-series forecasts of annual earnings. *Review of Accounting Studies* 17 (4): 944–968. https://doi.org/10.1007/s11142-012-9185-8.

Brown, L. D. and M. S. Rozeff. 1978. The superiority of analyst forecasts as measures of expectations: Evidence from earnings. *The Journal of Finance* 33 (1): 1–16. https://doi.org/10.1111/j.1540-6261.1978.tb03385.x.

Brown, L. D., A. C. Call, M. B. Clement, and N. Y. Sharp. 2015. Inside the "black box" of sell-side financial analysts. *Journal of Accounting Research* 53 (1): 1–47. https://doi.org/10.1111/1475-679x.12067.

Cao, S., W. Jiang, J. Wang, and B. Yang. 2024. From man vs. machine to man + machine: The art and ai of stock analyses. *Journal of Financial Economics* 160: 103910. https://doi.org/10.1016/j.jfineco.2024.103910.

Castelo, N., M. W. Bos, and D. R. Lehmann. 2019. Task-dependent algorithm aversion. *Journal of Marketing Research* 56 (5): 809–825. https://doi.org/10.1177/0022243719851788.

Chen, T., J. Harford, and C. Lin. 2015. Do analysts matter for governance? evidence from natural experiments. *Journal of Financial Economics* 115 (2): 383–410. https://doi.org/10.1016/j.jfineco.2014.10.002.

Chi, F., B.-H. Hwang, and Y. Zheng. 2025. The use and usefulness of big data in finance: Evidence from financial analysts. *Management Science* 71 (6): 4599–4621. https://doi.org/10.1287/mnsc.2022.02659.

Christ, M. H., M. J. Kim, and M. A. Yip. 2025. Word on the street: Field evidence of artificial intelligence use in sell-side equity research. *Working paper* . https://ssrn.com/abstract=5071836.

Clement, M. B. 1999. Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics* 27 (3): 285–303. https://doi.org/10.1016/s0165-4101(99)00013-0.

Clement, M. B. and S. Y. Tse. 2005. Financial analyst characteristics and herding behavior in forecasting. *The Journal of Finance* 60 (1): 307–341. https://doi.org/10.1111/j.1540-6261.2005.00731.x.

Cohen, L., D. Lou, and C. J. Malloy. 2020. Casting conference calls. *Management Science* 66 (11): 5015–5039. https://doi.org/10.1287/mnsc.2019.3423.

Coleman, B., K. J. Merkley, and J. Pacelli. 2022. Human versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations. *The Accounting Review* 97 (5): 221–244. https://doi.org/10.2308/tar-2020-0096.

de Silva, T., D. Thesmar, and S. Giglio. 2024. Noise in expectations: Evidence from analyst forecasts. *The Review of Financial Studies* 37 (5): 1494–1537. https://doi.org/10.1093/rfs/hhad091.

Deming, D. J. and L. B. Kahn. 2018. Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics* 36 (S1): S337–S369. https://doi.org/10.1086/694106.

Derrien, F. and A. Kecskés. 2013. The real effects of financial shocks: Evidence from exogenous changes in analyst coverage. *The Journal of Finance* 68 (4): 1407–1440. https://doi.org/10.1111/jofi.12042.

Devaraj, S. and R. Kohli. 2003. Performance impacts of information technology: Is actual usage the missing link? *Management Science* 49 (3): 273–289. https://doi.org/10.1287/mnsc.49.3.273.12736.

Dietvorst, B. J., J. P. Simmons, and C. Massey. 2015. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144 (1): 114–126. https://doi.org/10.1037/xge0000033.

DiMaggio, P. J. and W. W. Powell. 1983. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review* 48 (2): 147–160. https://doi.org/10.2307/2095101.

Domingos, P. 2012. A few useful things to know about machine learning. *Communications of the ACM* 55 (10): 78–87. https://doi.org/10.1145/2347736.2347755.

Driskill, M., M. P. Kirk, and J. W. Tucker. 2020. Concurrent earnings announcements and analysts' information production. *The Accounting Review* 95 (1): 165–189. https://doi.org/10.2308/accr-52489.

Frankel, R. and C. M. Lee. 1998. Accounting valuation, market expectation, and cross-sectional stock returns. *Journal of Accounting and Economics* 25 (3): 283–319. https://doi.org/10.1016/s0165-4101(98)00026-3.

Frankel, R., S. P. Kothari, and J. Weber. 2006. Determinants of the informativeness of analyst research. *Journal of Accounting and Economics* 41 (1–2): 29–54. https://doi.org/10.1016/j.jacceco.2005.10.004.

Fried, D. and D. Givoly. 1982. Financial analysts' forecasts of earnings: A better surrogate for market expectations. *Journal of Accounting and Economics* 4 (2): 85–107. https://doi.org/10.1016/0165-4101(82)90015-5.

Fügener, A., D. D. Walzner, and A. Gupta. 2026. Roles of artificial intelligence in collaboration with humans: Automation, augmentation, and the future of work. *Management Science* 72 (1): 538–557. https://doi.org/10.1287/mnsc.2024.05684.

Goldfarb, A., B. Taska, and F. Teodoridis. 2023. Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. *Research Policy* 52 (1): 104653. https://doi.org/10.1016/j.respol.2022.104653.

Google. 2022. Google COVID-19 community mobility reports. https://www.google.com/covid19/mobility/.

Green, T. C., R. Jame, S. Markov, and M. Subasi. 2014a. Access to management and the informativeness of analyst research. *Journal of Financial Economics* 114 (2): 239–255. https://doi.org/10.1016/j.jfineco.2014.07.003.

Green, T. C., R. Jame, S. Markov, and M. Subasi. 2014b. Broker-hosted investor conferences. *Journal of Accounting and Economics* 58 (1): 142–166. https://doi.org/10.1016/j.jacceco.2014.06.005.

Griliches, Z. 1957. Hybrid corn: An exploration in the economics of technological change. *Econometrica* 25 (4): 501–522. https://doi.org/10.2307/1905380.

Halevy, A., P. Norvig, and F. Pereira. 2009. The unreasonable effectiveness of data. *IEEE Intelligent Systems* 24 (2): 8–12. https://doi.org/10.1109/mis.2009.36.

Hall, B. H. and B. Khan. 2003. Adoption of new technology. Working Paper 9730, National Bureau of Economic Research. https://www.nber.org/papers/w9730.

Ham, C., R. N. Hann, M. Rabier, and W. Wang. 2024. Auditor skill demands and audit quality: Evidence from job postings. *Management Science* . https://doi.org/10.1287/mnsc.2022.01640.

Hirshleifer, D., S. S. Lim, and S. H. Teoh. 2009. Driven to distraction: Extraneous events and underreaction to earnings news. *The Journal of Finance* 64 (5): 2289–2325. https://doi.org/10.1111/j.1540-6261.2009.01501.x.

Hirshleifer, D., Y. Levi, B. Lourie, and S. H. Teoh. 2019. Decision fatigue and heuristic analyst forecasts. *Journal of Financial Economics* 133 (1): 83–98. https://doi.org/10.1016/j.jfineco.2019.01.005.

Hodder, L., P. E. Hopkins, and D. A. Wood. 2008. The effects of financial statement and informational complexity on analysts' cash flow forecasts. *The Accounting Review* 83 (4): 915–956. https://doi.org/10.2308/accr.2008.83.4.915.

Hong, H. and M. Kacperczyk. 2010. Competition and bias. *The Quarterly Journal of Economics* 125 (4): 1683–1725. https://doi.org/10.1162/qjec.2010.125.4.1683.

Hong, H., J. D. Kubik, and A. Solomon. 2000. Security analysts' career concerns and herding of earnings forecasts. *The RAND Journal of Economics* 31 (1): 121–144. https://doi.org/10.2307/2601032.

Huang, S. X., A. Hugon, T. Zhang, and W. Zheng. 2026. Generative ai and investment research: Evidence from analyst reports. *Working paper* . https://ssrn.com/abstract=6682498.

Jacob, J., T. Z. Lys, and M. A. Neale. 1999. Expertise in forecasting performance of security analysts. *Journal of Accounting and Economics* 28 (1): 51–82. https://doi.org/10.1016/s0165-4101(99)00016-6.

Jame, R., S. Markov, and M. C. Wolfe. 2022. Can fintech competition improve sell-side research quality? *The Accounting Review* 97 (4): 287–316. https://doi.org/10.2308/tar-2019-0266.

Jiao, Y. 2024. Managing decision fatigue: Evidence from analysts' earnings forecasts. *Journal of Accounting and Economics* 77 (1): 101615. https://doi.org/10.1016/j.jacceco.2023.101615.

Jordan, M. I. and T. M. Mitchell. 2015. Machine learning: Trends, perspectives, and prospects. *Science* 349 (6245): 255–260. https://doi.org/10.1126/science.aaa8415.

Kirk, M. P. and S. Markov. 2016. Come on over: Analyst/investor days as a disclosure medium. *The Accounting Review* 91 (6): 1725–1750. https://doi.org/10.2308/accr-51418.

Lehavy, R., F. Li, and K. J. Merkley. 2011. The effect of annual report readability on analyst following and the properties of their earnings forecasts. *The Accounting Review* 86 (3): 1087–1115. https://doi.org/10.2308/accr.00000043.

Li, F., R. Lundholm, and M. Minnis. 2013. A measure of competition based on 10-k filings. *Journal of Accounting Research* 51 (2): 399–436. https://doi.org/10.1111/j.1475-679X.2012.00472.x.

Lindebaum, D., M. Vesa, and F. den Hond. 2020. Insights from "the machine stops" to better understand rational assumptions in algorithmic decision making and its implications for organizations. *Academy of Management Review* 45 (1): 247–263. https://doi.org/10.5465/amr.2018.0181.

Mansfield, E. 1961. Technical change and the rate of imitation. *Econometrica* 29 (4): 741–766. https://doi.org/10.2307/1911817.

Marcus, G. and E. Davis. 2019. *Rebooting AI: Building Artificial Intelligence We Can Trust*. New York, NY, Pantheon Books.

Mayew, W. J. 2008. Evidence of management discrimination among analysts during earnings conference calls. *Journal of Accounting Research* 46 (3): 627–659. https://doi.org/10.1111/j.1475-679X.2008.00285.x.

Mayew, W. J., N. Y. Sharp, and M. Venkatachalam. 2013. Using earnings conference calls to identify analysts with superior private information. *Review of Accounting Studies* 18 (2): 386–413. https://doi.org/10.1007/s11142-012-9210-y.

McLaren, J. and S. Wang. 2020. Effects of reduced workplace presence on COVID-19 deaths: An instrumental-variables approach. NBER Working Paper 28275, National Bureau of Economic Research. https://www.nber.org/papers/w28275.

Merkley, K. J. 2014. Narrative disclosure and earnings performance: Evidence from r&d disclosures. *The Accounting Review* 89 (2): 725–757. https://doi.org/10.2308/accr-50649.

Merkley, K. J., R. Michaely, and J. Pacelli. 2017. Does the scope of the sell-side analyst industry matter? an examination of bias, accuracy, and information content of analyst reports. *The Journal of Finance* 72 (3): 1285–1334. https://doi.org/10.1111/jofi.12485.

Rogers, E. M. 2003. *Diffusion of Innovations*. New York, Free Press, 5th edition.

Shanthikumar, D. M. and I. S. Yoo. 2026. Beyond automation: AI and the human value of sell-side analysts (Working paper). https://ssrn.com/abstract=5040339.

Tambe, P. 2014. Big data investment, skills, and firm value. *Management Science* 60 (6): 1452–1469. https://doi.org/10.1287/mnsc.2014.1899.

Trott, S. and P. D. Rivière. 2024. Measuring and modifying the readability of English texts with GPT-4. In *Proceedings of the Third Workshop on Text Simplification, Accessibility and Readability (TSAR 2024)*, pages 126–134, Miami, FL. https://doi.org/10.18653/v1/2024.tsar-1.13.

van Binsbergen, J. H., X. Han, and A. Lopez-Lira. 2023. Man versus machine learning: The term structure of earnings expectations and conditional biases. *The Review of Financial Studies* 36 (6): 2361–2396. https://doi.org/10.1093/rfs/hhac085.

Welch, I. 2000. Herding among security analysts. *Journal of Financial Economics* 58 (3): 369–396. https://doi.org/10.1016/s0304-405x(00)00076-3.

Yezegel, A. 2023. The value of eliciting information: Evidence from sell-side analysts. *The Accounting Review* 98 (3): 459–486. https://doi.org/10.2308/TAR-2020-0308.

FIGURE 1. Trends in AI and Non-AI Technology Hiring

Panel A: AI Hiring

Panel B: Non-AI Technology Hiring

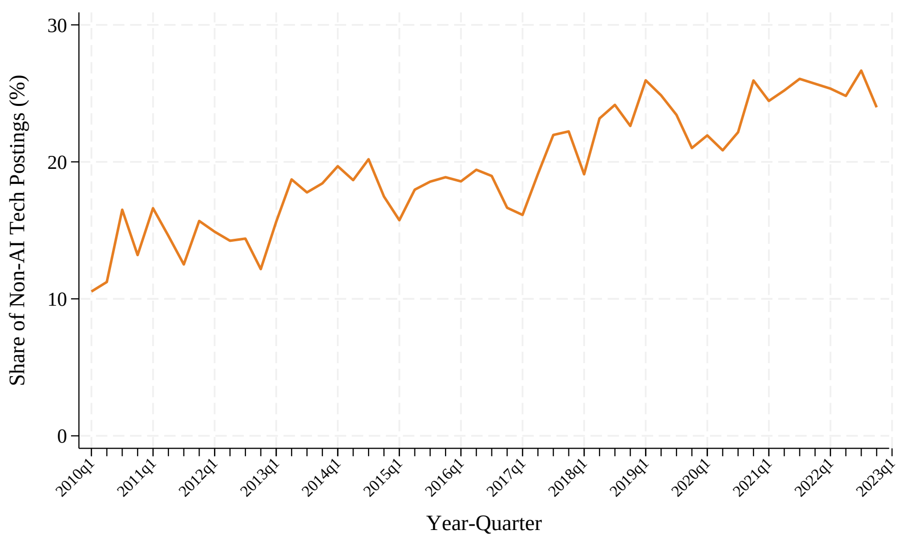

This figure plots the quarterly composition of technology-related hiring among research-related positions at brokerages from 2010 to 2022. **Panel A** plots the share of AI-related hiring (*AI Hiring*), computed at the brokerage-quarter level as a brokerage's AI-related postings as a percentage of its research-related postings, then averaged across the sample brokerages. **Panel B** plots the share of non-AI technology hiring (*Tech Hiring*), computed and averaged in the same way from postings requiring non-AI technology skills; AI-related postings are excluded, so the two panels capture mutually exclusive sets. The two series are distinct constructs rather than two labels for the same digitization trend: across the 1,436 active brokerage-quarters, their correlation is statistically indistinguishable from zero, both posting-weighted (r = -0.00, p = 0.88) and within brokerages (r = -0.03, p = 0.23).

FIGURE 2. How AI Enters Analyst Work and the Decision-Fatigue Predictions

Panel A: How AI Enters Analyst Work

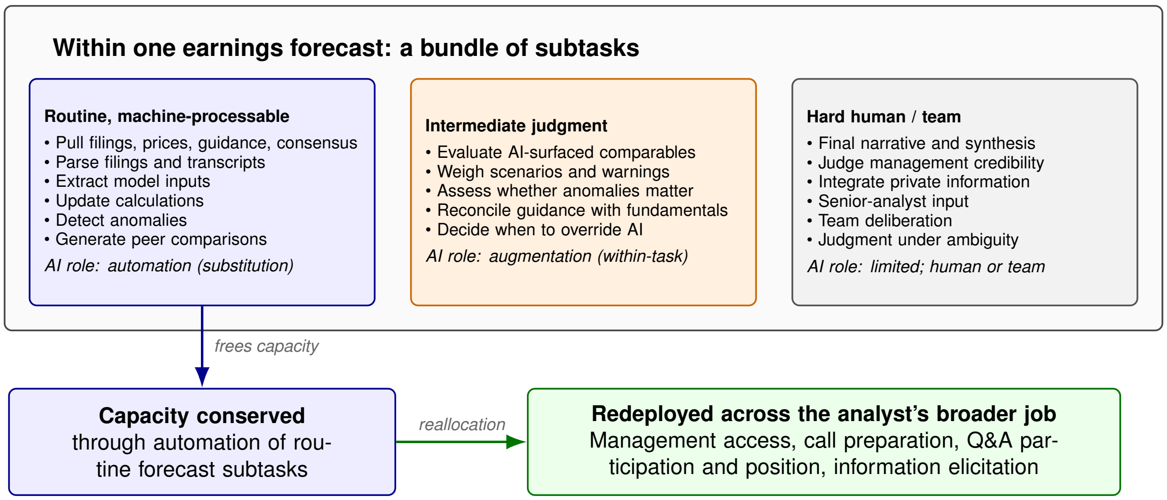

Panel B: Decision-Fatigue Predictions

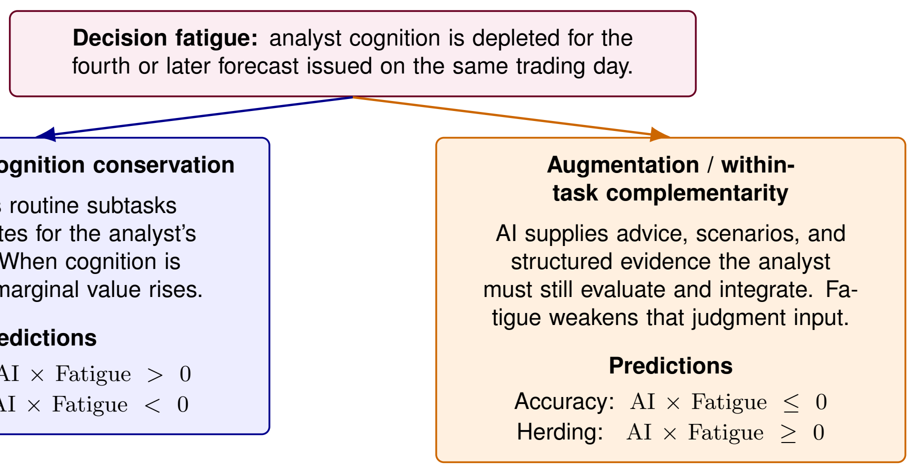

This figure presents our conceptual framework: how brokerage AI enters the production of an earnings forecast, how the capacity it conserves is redeployed, and how decision fatigue separates the automation and augmentation channels within the task. **Panel A** decomposes a single earnings forecast into a bundle of subtasks and identifies, for each group, the economic role AI plays: automation (substitution) on routine, machine-processable subtasks; within-task complementarity (augmentation) on intermediate judgment subtasks; and a limited direct role on the hard human or team subtasks. The capacity conserved by automating the routine subtasks can be redeployed across the analyst's broader job. **Panel B** translates this distinction into the decision-fatigue test: the automation channel and the within-task-complementarity channel imply opposite signs for the *AI Adoption* $\times$ fatigue interaction, so the estimated interaction discriminates between them.

TABLE 1. Sample Selection

Panel A: Brokerage Firm Sample

Procedure	**N**
I/B/E/S brokerages that issued EPS forecasts every year over 2018–2022	98
Deleting brokerage firms that cannot be identified in Lightcast database	(19)
Number of brokerage firms	79

Panel B: Earnings Forecast Sample

Procedure	**N**
Number of annual EPS forecasts in I/B/E/S	1,908,332
Number of annual EPS forecasts by analysts of listed brokers during 2012–2022	1,268,141
Deleting forecasts missing firm identifiers in Compustat or CRSP	(113,800)
Deleting forecasts by firms in financial industry and utility industry	(254,485)
Requiring at least 2 analysts issued forecasts in a given firm-year	(81,710)
Number of forecasts used in the main test	816,418

This table shows the sample selection process. **Panel A** presents the selection process of brokerage firm sample. **Panel B** presents the selection process for the annual EPS forecast sample.

TABLE 2. Summary Statistics

Panel A: Variables Used in the Analysis of the Determinants of AI Hiring

Variable	N	Mean	SD	Min	P25	P50	P75	Max
AI Hiring	2,916	0.005	1.011	-0.242	-0.242	-0.242	-0.242	15.472
Tech Hiring	2,916	-0.010	0.993	-0.675	-0.675	-0.675	0.902	2.478
Peer Adoption	2,916	0.019	1.003	-1.242	-0.733	-0.224	0.794	2.576
Broker Size	2,916	0.016	1.006	-0.738	-0.627	-0.378	0.287	5.082
Firm Coverage	2,916	0.043	0.968	-2.679	-0.677	0.156	0.825	1.522
Industry Coverage	2,916	0.037	0.979	-2.272	-0.659	0.111	0.888	1.464
Mean Experience	2,916	0.011	0.986	-4.413	-0.296	0.241	0.556	1.750
Inst Ownership	2,916	0.038	0.936	-4.080	-0.406	0.061	0.581	2.073
# of Institutions	2,916	0.021	0.989	-2.581	-0.642	0.118	0.732	2.129
Forecast Dispersion	2,916	0.004	1.011	-0.169	-0.168	-0.167	-0.165	8.039
Mean Coverage	2,916	0.031	0.961	-3.619	-0.425	0.109	0.628	2.617
Intangibles	2,916	0.021	0.971	-2.226	-0.423	0.155	0.531	2.855
Std Ret	2,916	-0.006	0.994	-1.381	-0.679	-0.272	0.340	4.195
MVE	2,916	0.016	0.992	-2.536	-0.618	0.082	0.664	2.258
Market to Book	2,916	0.000	1.000	-1.912	-0.613	-0.094	0.483	4.489

Panel B: Variables Used in the Main Tests

Variable	N	Mean	SD	Min	P25	P50	P75	Max
Relative Accuracy	816,418	-0.165	0.972	-3.031	-0.708	0.000	0.469	1.841
AI Adoption	816,418	0.230	0.317	0.000	0.000	0.000	0.411	1.000
AI Adoption $[t\text{-}8, t\text{-}1]$	816,418	0.260	0.336	0.000	0.000	0.039	0.497	1.000
AI Adoption $[t\text{-}12, t\text{-}1]$	744,928	0.278	0.347	0.000	0.000	0.062	0.541	1.000
AI Adoption $[t\text{-}16, t\text{-}1]$	674,162	0.285	0.349	0.000	0.000	0.081	0.574	1.000
Firm Experience	816,418	0.434	0.330	0.000	0.132	0.373	0.728	1.000
General Experience	816,418	0.438	0.344	0.000	0.129	0.359	0.779	1.000
Effort	816,418	0.606	0.312	0.000	0.389	0.600	0.889	1.000
Firms Followed	816,418	0.482	0.309	0.000	0.243	0.467	0.714	1.000
Forecast Age	816,418	0.550	0.315	0.000	0.272	0.536	0.804	1.000
Broker Size	816,418	0.366	0.316	0.000	0.104	0.297	0.485	1.000
Analyst Coverage	816,418	2.388	0.553	0.693	2.079	2.485	2.773	3.332
Lockdown	203,114	0.473	0.499	0.000	0.000	0.000	1.000	1.000

Panel C: Other Variables

Variable	N	Mean	SD	Min	P25	P50	P75	Max
<u>*Variables Used in Measure Validation*</u>								
$\ln(1+\textit{AI Patent Count})$	650	1.819	3.300	0.000	0.000	0.000	0.000	10.005
AI Patent Share	650	0.095	0.199	0.000	0.000	0.000	0.000	1.000
$\ln(\textit{Citation Weighted})$	650	0.797	1.617	0.000	0.000	0.000	0.000	6.682
$\ln(\textit{Cohort Adjusted})$	650	0.462	1.129	0.000	0.000	0.000	0.000	5.289
AI Disclosure	1,070	0.004	0.014	0.000	0.000	0.000	0.000	0.126

Table 2 continued from previous page

Variable	N	Mean	SD	Min	P25	P50	P75	Max
Variables Used in Cross-Sectional Tests								
Info Asymmetry	796,896	0.599	0.282	0.000	0.362	0.589	0.863	1.333
Disclosure Readability	737,317	25.544	1.339	16.354	24.606	25.514	26.418	37.452
Tech Background	366,204	0.194	0.395	0.000	0.000	0.000	0.000	1.000
Variables Used in the Mechanism Tests								
Fcast Under Fatigue	420,284	0.109	0.312	0.000	0.000	0.000	0.000	1.000
Herding	354,223	0.392	0.488	0.000	0.000	0.000	1.000	1.000
Participation	711,651	0.444	0.496	0.000	0.000	0.000	1.000	1.000
Rel Position	316,989	0.461	0.332	0.000	0.167	0.444	0.750	1.000
Question Position	316,989	4.762	3.364	1.000	2.000	4.000	7.000	27.000
Total Questioners	316,989	9.052	3.884	1.000	6.000	9.000	12.000	27.000
SRL	315,928	0.077	0.926	-2.890	-0.647	-0.018	0.738	3.831
Variables Used in the Capital- and Labor-Market Consequence Tests								
Informativeness	177,154	0.457	0.343	0.000	0.149	0.420	0.742	1.000
Hiring Share (Entry-level)	407	0.566	0.304	0.000	0.360	0.537	0.824	1.000
Hiring Share (Mid-level)	407	0.393	0.296	0.000	0.154	0.394	0.562	1.000
Hiring Share (Senior)	407	0.041	0.082	0.000	0.000	0.000	0.065	1.000

This table presents summary statistics. **Panel A** reports the standardized brokerage-quarter variables used in the analysis of the determinants of AI hiring (2,916 brokerage-quarters). **Panel B** reports the variables used in the main tests; the alternative-window measures *AI Adoption* $[t\text{-}8, t\text{-}1]$, *AI Adoption* $[t\text{-}12, t\text{-}1]$, and *AI Adoption* $[t\text{-}16, t\text{-}1]$ replace the baseline four-quarter window with the trailing eight, twelve, and sixteen quarters, respectively. **Panel C** reports other variables used in measure validation; the cross-sectional tests; the mechanism tests (the decision-fatigue, herding, and earnings-call management-access analyses); and the capital- and labor-market consequence tests (the price-informativeness analysis and the brokerage-year analysis of the experience composition of analyst hiring, 2012–2022). Variables used to partition the sample in the cross-sectional tests are reported in raw units; *Disclosure Readability* is the raw Gunning Fog Index, with higher values indicating less readable disclosures. All variables are defined in Appendix A.

TABLE 3. Validating the AI Adoption Measure

Panel A: Predicting AI Patenting from the AI Adoption Measure

	AI Patent Quantity				*AI Patent Quality*			
Dependent variable:	$\ln(1+\textit{AI Patent Count}_{j,t})$		*AI Patent Share*		$\ln(\textit{Citation Weighted}_{j,t})$		$\ln(\textit{Cohort Adjusted}_{j,t})$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI Adoption (standardized)	1.744***	0.632***	0.092***	0.030***	0.884***	0.314***	0.664***	0.343***
	(8.22)	(3.06)	(8.86)	(3.21)	(6.43)	(2.74)	(6.00)	(3.56)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	649	649	649	649	649	649	649	649
Adj. R^2	0.574	0.854	0.329	0.507	0.621	0.815	0.583	0.707
Fixed Effects	Yr-Qtr	Broker, Yr-Qtr	Yr-Qtr	Broker, Yr-Qtr	Yr-Qtr	Broker, Yr-Qtr	Yr-Qtr	Broker, Yr-Qtr

Panel B: Brokerages' AI Adoption and Disclosure about AI in Earnings Calls

Dependent variable:	*AI Disclosure*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.218***	0.224***	0.299***	0.258**	0.201**	0.192**
	(2.86)	(3.32)	(3.22)	(2.53)	(2.49)	(2.25)
Controls	No	Yes	No	Yes	No	Yes
Observations	950	950	950	950	950	950
Adj. R^2	0.058	0.074	0.145	0.174	0.210	0.231
Fixed Effects	None	None	Broker	Broker	Broker, Yr-Qtr	Broker, Yr-Qtr

This table validates the AI adoption measure (*AI Adoption*) against two independent measures of brokerage AI activity. **Panel A** relates *AI Adoption* (standardized) to brokerages' AI *patenting* (649 broker-quarters, 2012–2022): Columns 1–4 capture AI patent quantity, Columns 5–8 AI patent quality. **Panel B** relates AI discussion in brokerages' *own* earnings calls (*AI Disclosure*) to *AI Adoption* for the 24 publicly listed brokerages (2012–2022). Controls in both panels are *Broker Size*, *Firm Coverage*, *Industry Coverage*, *Mean Experience*, and *Mean Firm Experience*. All variables are defined in Appendix A. t-statistics based on standard errors clustered by brokerage are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 4. Determinants of Brokerages' AI Hiring versus Non-AI Technology Hiring

Dependent variable:	*AI Hiring*			*Tech Hiring* (non-AI)		
	(1)	(2)	(3)	(4)	(5)	(6)
Competitive environment						
Peer Adoption	0.192***	0.167***	0.175***			
	(5.43)	(4.73)	(5.31)			
Brokerage-level capacity						
Broker Size	0.117***	0.132***	0.132***	0.344***	0.335***	0.336***
	(3.55)	(4.12)	(4.16)	(26.79)	(23.55)	(24.71)
Firm Coverage	0.217***	0.207***	0.199***	0.204***	0.186***	0.184***
	(4.66)	(4.66)	(4.39)	(4.65)	(4.97)	(4.95)
Industry Coverage	-0.309***	-0.286***	-0.287***	0.065*	0.110***	0.108***
	(-6.19)	(-5.42)	(-5.57)	(1.72)	(3.14)	(3.14)
Analyst-level human capital						
Mean Experience	0.025**	0.023*	0.025*	-0.043***	-0.045***	-0.045***
	(2.04)	(1.79)	(1.98)	(-3.37)	(-3.35)	(-3.17)
Portfolio-level demand for analyst research						
Inst Ownership	0.032***	0.025**	0.024**	0.000	-0.011	-0.012
	(3.84)	(2.67)	(2.42)	(0.01)	(-0.97)	(-1.03)
# of Institutions	0.235***	0.198**	0.222***	0.022	-0.106*	-0.080
	(4.15)	(2.62)	(3.01)	(0.37)	(-1.76)	(-1.54)
Forecast Dispersion	-0.051***	-0.052***	-0.053***	0.002	0.012	0.012
	(-7.66)	(-7.34)	(-7.29)	(0.18)	(1.31)	(1.35)
Mean Coverage	-0.044	-0.002	0.005	-0.139***	-0.159***	-0.164***
	(-0.85)	(-0.04)	(0.09)	(-2.97)	(-6.38)	(-7.07)
Intangibles	-0.093***	-0.079***	-0.079***	-0.089***	-0.094***	-0.097***
	(-3.55)	(-2.88)	(-2.96)	(-4.04)	(-5.36)	(-5.66)
Std Ret	0.029	0.022	0.012	0.068***	0.114***	0.111***
	(1.52)	(0.94)	(0.51)	(4.38)	(4.28)	(3.97)
MVE	-0.112*	-0.123*	-0.158**	0.215***	0.367***	0.345***
	(-2.00)	(-1.99)	(-2.63)	(5.74)	(6.68)	(6.53)
Market to Book	-0.101***	-0.113***	-0.112***	-0.084***	-0.089***	-0.087***
	(-8.19)	(-8.93)	(-8.77)	(-5.64)	(-8.71)	(-8.51)
Observations	2,916	2,916	2,916	2,916	2,916	2,916
Adj. R^2	0.099	0.102	0.102	0.338	0.344	0.352
Fixed Effects	No	Year	Year×Qtr	No	Year	Year×Qtr

This table relates brokerage and portfolio characteristics to two hiring measures, estimating Equation (4). In Columns 1–3 the dependent variable is *AI Hiring*; in Columns 4–6 it is *Tech Hiring*, the brokerage's *non-AI* technology hiring, constructed exactly parallel to *AI Hiring*. *Peer Adoption* (peers' AI adoption) enters only the AI-hiring columns. The sample is the common set of broker-quarters over 2012–2022 for which both measures are available. All variables are standardized. Each block reports the no-fixed-effects, Year, and Year×Quarter specifications. All variables are defined in Appendix A. t-statistics based on standard errors clustered by year×quarter are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 5. Baseline Results: AI Adoption and Forecast Accuracy

Panel A: Baseline Estimates

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.067***	0.024***	0.140***	0.018**	0.192***	0.025***
	(9.99)	(3.26)	(14.99)	(2.36)	(16.01)	(2.97)
<u>*Control variables*</u>						
Firm Experience		0.038***		0.022***		0.046***
		(5.78)		(3.88)		(3.72)
General Experience		0.001		0.065***		0.068***
		(0.11)		(4.87)		(3.53)
Effort		-0.055***		-0.066***		-0.074***
		(-9.21)		(-12.39)		(-11.94)
Broker Size		-0.023***		-0.009		0.025*
		(-2.69)		(-1.00)		(1.85)
Forecast Age		-1.403***		-1.412***		-1.432***
		(-176.29)		(-178.70)		(-176.18)
Firms Followed		0.019**		0.001		0.001
		(2.46)		(0.15)		(0.13)
Analyst Coverage		0.005		0.005		0.032***
		(1.22)		(1.30)		(3.78)
Observations	816,418	816,418	816,244	816,244	813,205	813,205
Adj. R^2	0.001	0.207	0.015	0.223	0.051	0.269
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm

Panel B: Alternative Temporal Windows to Capture Brokerages' AI Adoption

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption [t-8, t-1]	0.016**	0.024***				
	(1.96)	(2.67)				
AI Adoption [t-12, t-1]			0.014*	0.027***		
			(1.69)	(2.76)		
AI Adoption [t-16, t-1]					0.028***	0.042***
					(3.07)	(3.78)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	816,244	813,205	744,764	741,847	674,011	671,152
Adj. R^2	0.223	0.269	0.225	0.272	0.225	0.273
Fixed Effects	Analyst	Analyst×firm	Analyst	Analyst×firm	Analyst	Analyst×firm

Panel C: AI Adoption Measured Based on Stock AI Hirings

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption (Stock)	0.068***	0.015**	0.226***	0.021**	0.391***	0.032***
	(9.56)	(2.02)	(17.98)	(2.27)	(20.22)	(2.69)
Controls	No	Yes	No	Yes	No	Yes
Observations	647,461	647,461	647,331	647,331	645,157	645,157
Adj. R^2	0.001	0.208	0.014	0.221	0.049	0.265
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm

This table presents the results from estimating Equation (5). The sample includes all annual analyst EPS forecasts over 2012–2022. **Panel A** reports the baseline estimates. **Panel B** measures *AI Adoption* over longer trailing windows of the flow measure: eight quarters ($[t\text{-}8, t\text{-}1]$), twelve quarters ($[t\text{-}12, t\text{-}1]$), and sixteen quarters ($[t\text{-}16, t\text{-}1]$), with the control variables of Panel A included throughout. **Panel C** replaces the flow measure with a *stock* measure: *AI Adoption* (Stock) is the average share of AI-related hiring in the brokerage's total hiring from 2010Q1 to quarter t-1, standardized within each firm-year; the sample retains only brokerages with complete coverage in the Lightcast job-posting data from 2010Q1 onward, so that every brokerage enters the stock measure with the same length of hiring history. All variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 6. AI Adoption and Forecast Accuracy During Pandemic Workplace Lockdowns: Alternative AI Measurement Windows

Dependent variable:	*Relative Accuracy*							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Office Lockdown:	No	Yes	No	Yes	No	Yes	No	Yes
AI Adoption [t-4, t-1]	0.231***	0.047						
	(6.39)	(1.44)						
AI Adoption [t-8, t-1]			0.160***	0.000				
			(3.58)	(0.00)				
AI Adoption [t-12, t-1]					0.321***	0.018		
					(5.64)	(0.28)		
AI Adoption [t-16, t-1]							0.393***	0.107
							(6.77)	(1.61)
χ^2 Stat of Diff	13.07***		5.87**		13.06***		10.61***	
(p-value)	(<0.001)		(0.016)		(<0.001)		(0.001)	
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	98,104	89,875	98,104	89,875	98,104	89,875	98,104	89,875
Adj. R^2	0.300	0.375	0.299	0.375	0.300	0.375	0.300	0.375
Fixed Effects	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm

This table examines whether the forecast-accuracy benefit of brokerage AI adoption during the COVID-19 pandemic depends on analysts' physical workplace access, across alternative trailing windows for measuring *AI Adoption*; each column re-estimates the baseline analyst-firm fixed-effects specification (Equation 5; Table 5, Column 6). The sample comprises analyst-firm forecasts over 2020–2022. The columns split the sample by *Lockdown*, set to *Yes* if workplace visits in the brokerage's headquarters county fell below 60% of the January 2020 baseline in the Google COVID-19 Community Mobility Reports, and *No* otherwise; *AI Adoption* is estimated separately in each subsample, and the χ^2 statistic tests equality of its coefficient across the two. The control variables are those in Table 5. All variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 7. **AI Adoption and Forecast Accuracy: Firms' Information Environment and Analyst Technical Background**

Dependent variable:	*Relative Accuracy*							
Measure:	*Forecast Timing*		*Info Asymmetry*		*Disclosure Readability*		*Analyst Background*	
Split:	Early	Later	Low	High	High	Low	Technical	Non-technical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI Adoption	0.046***	-0.001	0.030**	0.011	0.045***	0.007	0.070**	0.010
	(3.39)	(-0.12)	(2.47)	(0.90)	(5.18)	(0.79)	(2.12)	(0.70)
χ^2 Stat of Diff	33.78***		18.05***		5.64**		2.73*	
(p-value)	(<0.001)		(<0.001)		(0.018)		(0.098)	
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	387,876	418,346	394,724	393,993	363,042	366,279	66,628	274,774
Adj. R^2	0.134	0.206	0.380	0.208	0.279	0.300	0.252	0.273
Fixed Effects	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm

This table presents the results from estimating Equation (5), partitioning the sample along cross-sectional dimensions of the information environment and analyst background. The sample includes all annual analyst EPS forecasts over 2012–2022. The control variables are those in Table 5. All variables are defined in Appendix A. The sample is partitioned by forecast timing; by information asymmetry among analysts (following Barron et al. (1998)); by disclosure readability (measured as the reverse-coded Gunning Fog Index of the firm's disclosures so that higher values denote more readable disclosures, with the High and Low columns splitting the sample at the median); and by analyst technical background (Technical versus Non-technical, i.e., with versus without a STEM education identified from LinkedIn/Revelio profiles; 1,879 matched analysts). The disclosure-readability results are robust to measuring readability with the SMOG and Automated Readability indices in place of the Fog Index. For each partition, the χ^2 statistic tests the difference in the *AI Adoption* coefficient across subsamples. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 8. Mitigating The Adverse Effects of Decision Fatigue

Panel A: AI Adoption and Decision Fatigue

Dependent variable:	*Relative Accuracy*		
	(1)	(2)	(3)
Fcast Under Fatigue	-0.046***	-0.033***	-0.037***
	(-9.68)	(-7.95)	(-8.81)
AI Adoption	0.015	0.014	0.021
	(1.40)	(1.24)	(1.59)
Fcast Under Fatigue $\times$ *AI Adoption*	0.030***	0.026**	0.022**
	(2.69)	(2.47)	(2.03)
Controls	Yes	Yes	Yes
Observations	420,284	420,016	413,665
Adj. R^2	0.205	0.223	0.276
Fixed Effects	No	Analyst	Analyst$\times$firm

Panel B: AI Adoption and Analysts' Herding Behavior Under Decision Fatigue

Dependent variable:	*Herding* (0/1)			
	(1)	(2)	(3)	(4)
Fcast Under Fatigue	0.316***	0.319***	0.264***	0.263***
	(23.17)	(23.24)	(14.78)	(14.70)
AI Adoption	-0.009	0.008	-0.098***	-0.006
	(-0.74)	(0.60)	(-4.17)	(-0.25)
Fcast Under Fatigue $\times$ *AI Adoption*	-0.111***	-0.114***	-0.182***	-0.178***
	(-3.26)	(-3.34)	(-4.43)	(-4.32)
Controls	No	Yes	No	Yes
Observations	354,223	354,223	323,508	323,508
Pseudo R^2	0.002	0.004	0.001	0.005
Fixed Effects	No	No	Analyst$\times$firm	Analyst$\times$firm

This table presents the results from estimating Equations (5) and (7). The sample includes all annual EPS forecasts over 2012–2022. The control variables are those in Table 5. **Panel A** adds *Fcast Under Fatigue*, a dummy equal to 1 if the forecast is the fourth or later issued by the same analyst on the same day, and its interaction with *AI Adoption*; following Jiao (2024), only forecasts made during 9:00 to 23:00 are included. **Panel B** estimates Equation (7). All variables are defined in Appendix A. t-statistics (Panel A) and z-statistics (Panel B) based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 9. Brokerages' AI Adoption and Analyst Engagement and Information Elicitation in Earnings Calls

Panel A: Engagement in Earnings Calls

Dependent variable:	*Participation*		*Rel Position*		*Question Position*	
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.172**	0.175***	-0.141**	-0.119**	-1.518***	-1.297**
	(2.34)	(2.58)	(-2.35)	(-2.00)	(-2.69)	(-2.33)
Controls	No	Yes	No	Yes	No	Yes
Firm×quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Analyst×firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	647,791	647,791	286,502	286,502	286,502	286,502
Adj. R^2	0.351	0.380	0.290	0.291	0.457	0.458

Panel B: Information Elicitation from Management in Earnings Calls

Dependent variable:	*SRL*		Ln(*SRL*)	
	(1)	(2)	(3)	(4)
AI Adoption	0.343**	0.339*	0.366**	0.356**
	(1.99)	(1.96)	(2.24)	(2.17)
Controls	No	Yes	No	Yes
Firm×quarter FE	Yes	Yes	Yes	Yes
Analyst×firm FE	Yes	Yes	Yes	Yes
Observations	286,495	286,495	286,495	286,495
Adj. R^2	0.037	0.037	0.041	0.041

This table relates brokerage AI adoption to analysts' behavior in earnings calls. The unit of observation is the analyst-firm-quarter; the sample comprises all analysts covering the firm, identified from S&P Capital IQ transcripts over 2012–2022. **Panel A** measures *engagement* (*Participation*, *Rel Position*, *Question Position*); because *Rel Position* and *Question Position* (Columns 3–6) are defined only for analysts who ask at least one question, those columns are estimated on the subsample of participating analysts. **Panel B** measures *information elicitation* using the standardized response length (*SRL*) measure of Yezegel (2023). The control variables are those in Table 5. All variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE 10. Capital- and Labor-Market Consequences of Brokerages' AI Adoption

Panel A: Analyst Research Informativeness (Capital Markets)

Dependent variable:	*Informativeness*			
	(1)	(2)	(3)	(4)
AI Adoption	0.018***	0.008*	0.009*	0.002
	(4.13)	(1.82)	(1.87)	(0.36)
AI Disclosure			-0.010*	-0.010*
			(-1.73)	(-1.74)
AI Adoption × *AI Disclosure*			0.011**	0.009*
			(2.17)	(1.87)
Controls	No	Yes	No	Yes
Observations	166,399	166,399	107,813	107,813
Adj. R^2	0.165	0.167	0.151	0.152
Fixed Effects	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm

Panel B: Experience Composition of Next-Year Analyst Hiring (Labor Markets)

	Share of year-(t+1) analyst postings		
	Entry-level (1)	Mid-level (2)	Senior (3)
AI Adoption	-0.121***	0.077***	0.226***
	(-4.01)	(2.76)	(4.49)
Determinant controls	Yes	Yes	Yes
Within-year standardized	Yes	Yes	Yes
Observations	370	370	370
R^2	0.199	0.190	0.204

This table reports two consequences of brokerages' AI adoption. **Panel A** relates analyst research informativeness to AI adoption and its public disclosure. *Informativeness* is the Frankel et al. (2006) measure based on market-adjusted abnormal returns at the analyst-firm-year level (robust to size- and industry-adjusted models); Columns 1–2 relate *AI Adoption* to *Informativeness* (Equation 9) on the full sample, and Columns 3–4 add *AI Disclosure*, the brokerage's standardized AI-discussion intensity in its own earnings calls, and its interaction with *AI Adoption* on the subsample of publicly listed brokerages. The control variables are those in Table 5. **Panel B** relates the experience composition of a brokerage's next-year analyst hiring to its AI adoption, using an annual brokerage panel for 2012–2022. A posting is classified as entry-level if it requires two or fewer years of experience, is an internship, or carries a junior title, as senior if it carries a senior, vice-president, or director title, and as mid-level otherwise; the dependent variable is the share of year-(t+1) analyst postings in each category (Equation 10). Following the determinants analysis (Table 4), every variable is standardized within year and the specification includes the same brokerage- and portfolio-level controls and pools across brokerages. In levels, AI-adopting brokerages post more analyst openings at every experience level (untabulated), so Panel B reflects a compositional tilt within an expanding hiring program. All variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst in Panel A and by brokerage in Panel B are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

APPENDIX A

Variable Definitions

Variable	Definition
1. Variables Used in the Analysis of the Determinants of AI Hiring	
AI Hiring$_{b,t}$	The standardized share of full-time research-related job postings at brokerage *b* in quarter *t* requiring at least one AI-related skill. (Source: Lightcast)
Tech Hiring$_{b,t}$	The number of postings requiring non-AI technology skills as a share of total research-related postings at brokerage *b* in quarter *t*, excluding AI-related postings so that this and *AI Hiring* capture mutually exclusive sets. (Source: Lightcast)
Peer Adoption$_{b,t}$	The standardized number of brokerages in the same size quantile as brokerage *b* that have made at least one AI-related hire up to the beginning of quarter *t*. (Source: Lightcast)
Broker Size$_{b,t}$	The standardized number of analysts employed by brokerage *b* in quarter *t*, measured at the brokerage-quarter level (distinct from the analyst-firm-year, range-scaled *Broker Size* used as a control in the main tests). (Source: I/B/E/S)
Firm Coverage$_{b,t}$	The standardized number of unique firms covered by analysts affiliated with brokerage *b* in quarter *t*. (Source: I/B/E/S)
Industry Coverage$_{b,t}$	The standardized number of unique SIC industries of firms covered by analysts affiliated with brokerage *b* in quarter *t*. (Source: I/B/E/S, Compustat)
Mean Experience$_{b,t}$	The standardized average experience of analysts affiliated with brokerage *b* in quarter *t*, experience being the number of days since their first forecast. (Source: I/B/E/S)
Inst Ownership$_{b,t}$	The standardized average institutional ownership across all firms followed by brokerage *b* in quarter *t*. (Source: I/B/E/S, LSEG 13F)
# of Institutions$_{b,t}$	The standardized average number of institutional holders across all firms covered by brokerage *b* in quarter *t*. (Source: I/B/E/S, LSEG 13F)
Forecast Dispersion$_{b,t}$	The standardized average forecast dispersion, the standard deviation of analyst earnings forecasts, across firms in brokerage *b*'s portfolio in quarter *t*. (Source: I/B/E/S)
Mean Coverage$_{b,t}$	The standardized average number of analysts covering each firm in brokerage *b*'s portfolio in quarter *t*. (Source: I/B/E/S)
Intangibles$_{b,t}$	The standardized average ratio of intangible assets to total assets across firms followed by brokerage *b* in quarter *t*. (Source: I/B/E/S, Compustat)
Std Ret$_{b,t}$	The standardized average standard deviation of daily stock returns over the past year across firms followed by brokerage *b* in quarter *t*. (Source: I/B/E/S, CRSP)
MVE$_{b,t}$	The standardized average market capitalization of firms covered by brokerage *b* in quarter *t*. (Source: I/B/E/S, CRSP)
Market to Book$_{b,t}$	The standardized average market-to-book ratio of firms followed by brokerage *b* in quarter *t*. (Source: I/B/E/S, Compustat)
2. Variables Used in the Main Tests	
Relative Accuracy$_{i,j,t}$	Analyst *i*'s EPS forecast error for firm *j* at time *t* minus the median forecast error across all analysts covering firm *j* in year *t*, divided by the standard deviation of those analysts' forecast errors. (Source: I/B/E/S)
AI Adoption$_{i,j,t}$	AI investment, proxied by the share of AI-related job postings over a trailing window, then standardized as analyst *i*'s brokerage AI-investment measure at time *t* minus the minimum among brokerages of analysts covering firm *j* in year *t*, scaled by the range. The baseline uses the trailing four-quarter window [*t*-4, *t*-1]; the variants [*t*-8, *t*-1], [*t*-12, *t*-1], and [*t*-16, *t*-1] use the trailing eight, twelve, and sixteen quarters. (Source: Lightcast)

Appendix A continued from previous page

Variable	**Definition**
AI Adoption$_{i,j,t}$ (Stock)	A *stock* measure of brokerage AI adoption: the average share of AI-related hiring in the brokerage's total hiring from 2010Q1 to quarter t-1, min-max scaled within firm-year, for brokerages present in the job-posting sample since 2010Q1. (Source: Lightcast)
Firm Experience$_{i,j,t}$	Analyst i's days of firm-specific experience covering firm j at time t minus the minimum among analysts covering firm j, scaled by the range across those analysts in year t. (Source: I/B/E/S)
General Experience$_{i,j,t}$	Analyst i's days of general experience at time t minus the minimum among analysts covering firm j, scaled by the range across those analysts in year t. (Source: I/B/E/S)
Effort$_{i,j,t}$	The number of forecasts analyst i issues for firm j in year t minus the minimum among analysts covering firm j, scaled by the range across those analysts in year t. (Source: I/B/E/S)
Firms Followed$_{i,j,t}$	The number of firms analyst i follows in year t minus the minimum among analysts covering firm j, scaled by the range across those analysts in year t. (Source: I/B/E/S)
Forecast Age$_{i,j,t}$	The number of days from analyst i's forecast date to the earnings announcement in year t minus the minimum among analysts covering firm j, scaled by the range across those analysts in year t. (Source: I/B/E/S)
Broker Size$_{i,j,t}$	The number of analysts employed by analyst i's brokerage in year t minus the minimum among analysts covering firm j, scaled by the range across those analysts in year t. (Source: I/B/E/S)
Analyst Coverage$_{i,j,t}$	The logarithm of the number of analysts who cover firm j in year t. (Source: I/B/E/S)
Lockdown$_{i,j,t}$	A dummy equal to one if, in the quarter of analyst i's forecast, workplace visits in the county of the brokerage's headquarters fell below 60% of the January 2020 pre-pandemic baseline, and zero otherwise. (Source: Google COVID-19 Community Mobility Reports)

3. Other Variables

Variables Used in Measure Validation	
AI Patent Count$_{b,t}$	The number of AI-related patents filed by brokerage b in quarter t, from USPTO filings matched to the brokerage. (Source: USPTO, PatentsView)
AI Patent Share$_{b,t}$	Brokerage b's AI patents divided by its total patents filed in quarter t, bounded between zero and one. (Source: USPTO, PatentsView)
Citation Weighted$_{b,t}$	The sum of $(1+\text{forward citations})$ over brokerage b's AI patents in quarter t. (Source: USPTO, PatentsView)
Cohort Adjusted$_{b,t}$	The sum over brokerage b's AI patents in quarter t of each patent's forward citations scaled by the average for patents filed in the same year. (Source: USPTO, PatentsView)
AI Disclosure$_{b,t}$	The share (in percent) of AI-related words in all management speech (prepared remarks and Q&A, across all transcript types) in brokerage b's *own* earnings calls and investor communications in quarter t. AI terms comprise *artificial intelligence*, *machine learning*, *generative AI*, *deep learning*, *large language model*, *neural network*, *natural language processing*, and *AI*. Defined for the publicly listed brokerages in the sample. (Source: S&P Capital IQ)
Mean Firm Experience$_{b,t}$	The average firm-specific experience (log number of days covering portfolio firms) of analysts employed at brokerage b in quarter t. (Source: I/B/E/S)
Variables Used in Cross-Sectional Tests	
Forecast Timing$_{i,j,t}$	The forecast horizon: the number of days from analyst i's forecast for firm j in year t to the earnings-announcement date, so larger values denote earlier forecasts. The cross-sectional test partitions forecasts at the median. (Source: I/B/E/S)

Variable	**Definition**
Info Asymmetry$_{i,j,t}$	Information asymmetry among analysts covering firm j in year t, based on forecast dispersion following Barron et al. (1998). (Source: I/B/E/S)
Disclosure Readability$_{i,j,t}$	Readability of firm j's public disclosures. Table 2 reports the raw Gunning Fog Index ($0.4 \times (\text{words}/\text{sentences}) + 40 \times (\text{complex words}/\text{words})$), for which higher values denote less readable (more linguistically complex) disclosures. The cross-sectional test in Table 7 reverse-codes the index and splits at the median, so that the High-readability subsample comprises forecasts on firms with below-median Fog. (Source: SEC Analytics Suite)
Tech Background$_{i,j,t}$	A dummy equal to one if analyst i has a STEM educational background, identified from LinkedIn/Revelio profiles, and zero otherwise. (Source: Revelio Labs)
<u>*Variables Used in the Mechanism Tests*</u>	
Fcast Under Fatigue$_{i,j,t}$	A dummy equal to one if the forecast is the fourth or later forecast issued by analyst i on the same day (issued between 9:00 and 23:00), and zero otherwise. (Source: I/B/E/S)
Herding$_{i,j,t}$	A dummy equal to one if the forecast lies between analyst i's own prior forecast and the consensus, and zero otherwise. (Source: I/B/E/S)
Participation$_{i,j,t}$	A dummy equal to one if analyst i asks at least one question on firm j's earnings call in quarter t, and zero otherwise; the sample comprises all analysts covering firm j. (Source: S&P Capital IQ)
Rel Position$_{i,j,t}$	Analyst i's normalized place in the question queue of firm j's earnings call in quarter t, $(k\text{-}1)/(K\text{-}1)$, where k is the order in which analyst i first asks and K is the number of questioners; it ranges from 0 (first) to 1 (last), with smaller values denoting an earlier position. (Source: S&P Capital IQ)
Question Position$_{i,j,t}$	The raw rank k at which analyst i first asks a question on firm j's earnings call in quarter t (1 = first questioner). (Source: S&P Capital IQ)
Total Questioners$_{i,j,t}$	The number of analysts K asking at least one question on firm j's earnings call in quarter t. (Source: S&P Capital IQ)
SRL$_{i,j,t}$	Standardized response length: the total words management uses to answer analyst i's questions on firm j's earnings call in quarter t (each answer attributed to the most recently questioning analyst), standardized within the call across questioning analysts to mean zero and unit standard deviation, so higher values indicate longer responses to analyst i than to her peers. (Source: S&P Capital IQ)
<u>*Variables Used in the Capital- and Labor-Market Consequence Tests*</u>	
Informativeness$_{i,j,t}$	Analyst i's *Analyst Informativeness* for firm j in year t minus the minimum among analysts covering firm j in year t, scaled by the range. *Analyst Informativeness* is the sum of absolute daily abnormal returns on analyst i's forecast days for firm j in year t, divided by the sum of daily abnormal returns within year t, and further divided by the number of forecasts analyst i makes on firm j in year t. Abnormal returns are adjusted for market, size, and industry. (Source: I/B/E/S, CRSP)
Hiring Share$^{g}_{b,t}$	The share of brokerage b's analyst research-related job postings in experience group g, where g is entry-level, mid-level, or senior, in year t. A posting is classified as entry-level if it requires two or fewer years of experience, is an internship, or carries a junior title; as senior if it carries a senior, vice-president, or director title; and as mid-level otherwise. (Source: Lightcast)

ONLINE APPENDIX

"AI-Powered Analysts"

This Online Appendix supplements "AI-Powered Analysts." Section OA I reports additional tests: the relation between AI adoption and the analyst experience gap, forecast characteristics, forecast optimism, alternative fixed-effects structures, and a brokerage-level version of the accuracy test. Section OA II collects robustness checks of the baseline forecast-accuracy estimates, together with the binned-scatter figures for the validation and baseline results. Section OA III re-examines the decision-fatigue results under an alternative error metric and reports the lag structure of the baseline relation.

OA I. ADDITIONAL TESTS

Does AI Adoption Level the Playing Field?

Whether new technology levels the playing field between less- and more-skilled workers is a long-running question, and analyst experience offers a natural setting: forecasting performance rises with experience (Clement 1999; Jacob et al. 1999), as seasoned analysts accumulate forecasting heuristics, firm-specific knowledge, and access to management, while brokerage AI automates the processing of structured public data and may substitute for some of the advantages that experience confers. We examine this by interacting *AI Adoption* with indicators for below-median firm-specific and general experience (*Low Firm Experience* and *Low General Experience*), where firm-specific experience is the analyst's tenure covering the firm and general experience is time since the analyst's first appearance in I/B/E/S.[42]

Table OA2, Panel A, reports the estimates. Both interactions are positive and significant across specifications, while the main effect of *AI Adoption*, the gain for more-experienced analysts, is small and insignificant, a pattern consistent with the benefit of brokerage AI concentrating among less-experienced analysts. We interpret this result with caution, however. Experience is highly

[42] Because experience varies across analysts covering the same firm-year, unlike the firm-level partitions of Section V, it survives the within-firm-year standardization and can enter as an interaction with *AI Adoption* rather than through a sample split.

correlated with age, and younger analysts are more likely to use AI tools in the first place, so the estimated interactions may reflect differences in AI usage across experience levels rather than a differential benefit of AI at the same level of usage. Because our measure captures the brokerage's AI capability rather than the analyst's usage of it, we cannot separate the two; we therefore regard the evidence as suggestive and do not argue that AI levels the playing field.

Forecast Characteristics Other Than Accuracy

Our primary focus is the *quality* of analysts' forecasts; as an additional test, we ask whether brokerage AI also raises their *productivity*. We examine three forecast characteristics, each measured at the analyst-firm-quarter level: *Frequency*, the number of EPS forecasts the analyst issues for the firm in the quarter; *Thoroughness*, the number of distinct forecast horizons the analyst provides; and *Timeliness*, an indicator equal to one if the analyst's first forecast revision following the firm's earnings announcement arrives within two days. *Frequency* and *Thoroughness* enter as within-firm-quarter ranks scaled to the unit interval, and *Timeliness* is a binary indicator. We relate each to *AI Adoption* under the baseline analyst-firm fixed-effects specification, estimating *Frequency* and *Thoroughness* by OLS and the binary *Timeliness* by conditional logit.

Table OA2, Panel B, reports the results. *AI Adoption* is positive and significant for all three characteristics. Analysts at more AI-intensive brokerages forecast more often, cover more forecast horizons, and revise more quickly after earnings news. Brokerage AI thus increases the quantity, breadth, and speed of analysts' output, consistent with AI easing the processing constraints that otherwise limit how much research an analyst can produce.

Forecast Optimism

A natural alternative explanation is that AI improves accuracy by curbing optimistic bias, imposing a more data-driven discipline that limits reliance on subjective judgment. Forecasting, however, remains a human task. Analysts still choose how to weigh model output and which information to emphasize, and the strategic incentives behind optimism, such as preserving management access and client relationships, are not addressed by algorithmic tools. AI need not, therefore, reduce optimism. We test this using relative forecast optimism as the dependent variable:

$$Relative\ Optimism_{i,j,t} = \frac{OPT_{i,j,t} - \text{Median}\left(OPT_{\cdot,j,t}\right)}{\text{SD}\left(OPT_{\cdot,j,t}\right)}, \tag{OA.1}$$

where $OPT_{i,j,t}$ is analyst i's forecast optimism for firm j in fiscal year t, the signed forecast error (forecast minus realized EPS) scaled by the absolute value of realized EPS, and the median and standard deviation are taken across all analysts covering firm j in fiscal year t. Higher values denote more optimistic forecasts relative to peers.

Table OA2, Panel C, presents the estimates. *AI Adoption* shows no significant negative association with *Relative Optimism*, so the accuracy gains we document do not stem from a reduction in optimistic bias.

Alternative Fixed Effects

Finally, we confirm that the baseline result does not depend on a particular treatment of unobserved heterogeneity. We re-estimate Equation (5) under three alternative fixed-effects structures: broker, broker-firm, and analyst-firm-year.

Table OA2, Panel D, reports the results. The association between *AI Adoption* and *Relative Accuracy* is positive and significant across Columns 1–5, spanning broker, broker-firm, and analyst-firm-year fixed effects. In the most saturated specification, analyst-firm-year fixed effects with controls (Column 6), the coefficient remains positive but loses significance, as expected. AI hiring is stable at the broker level over short horizons, so analyst-firm-year effects absorb most of the variation in *AI Adoption* and leave little for identification. The baseline finding is otherwise robust across these structures.

Accuracy at the Brokerage Level

Because the treatment varies at the brokerage level, a natural question is whether the accuracy relation is visible when the analysis is conducted at that level outright. To examine this, we use *Average Relative AFE* as the dependent variable, where *AFE* is the absolute forecast error scaled by the beginning-of-year stock price. For each forecast, the relative *AFE* subtracts from the price-scaled error the mean price-scaled error across all analysts' forecasts for the same firm in the same quarter, and *Average Relative AFE* averages these relative errors across a brokerage's forecasts in each quarter. The within-firm-quarter adjustment cancels variation that originates in differences across coverage portfolios, since brokerages cover different sets of firms whose earnings are differentially difficult to forecast, and renders the averages comparable across brokerages. In line with the brokerage-level outcome, the explanatory variables, including *AI Adoption*, are all measured at the brokerage level. *AI Adoption* is the brokerage's trailing four-quarter AI-hiring share, and the controls are the brokerage-level determinants of Section IV other than *Peer Adoption*, with the regressors non-scaled (i.e., not min-max standardized).

Table OA2, Panel E reports the estimates on the sample of 2,916 brokerage-quarters, identical to the estimation sample of the determinants analysis, with and without the controls and with quarter fixed effects added in stages. *AI Adoption* enters negatively in every column and is statistically significant throughout, indicating that brokerages that adopt more AI produce forecasts with errors further below the consensus error on the same firms in the same periods. The positive relation between brokerage AI adoption and forecast accuracy therefore holds at the brokerage level as well.

OA II. ROBUSTNESS OF THE BASELINE ESTIMATES

An Alternative Measure of Forecast Error

The baseline relation is also robust to replacing the Clement-style relative-accuracy measure with a direct firm-year fixed-effects specification. Table OA3 re-estimates Equation (5) using absolute forecast error (AFE), scaled by beginning-of-year price and winsorized, as the dependent variable. The regressors are not min-max standardized within firm-year, and every specification includes firm-year fixed effects; the subsequent specifications pair these fixed effects with analyst fixed effects or analyst-firm fixed effects. AI Adoption is negative and statistically significant throughout, indicating that analysts at more AI-intensive brokerages make smaller absolute forecast errors. Thus, the baseline inference does not depend on the construction of Relative Accuracy or the within-firm-year scaling of the regressors. Table OA8 applies the same approach to the decision-fatigue analysis. The interaction between AI Adoption and Fcast Under Fatigue is negative, as expected when lower AFE denotes greater accuracy, and remains statistically significant in the controlled specifications with firm-year fixed effects, firm-year and analyst fixed effects, and firm-year and analyst-firm fixed effects.

Continuously Active Brokerages

The brokerage panel is unbalanced—brokerages enter and exit I/B/E/S over the sample period, and entry and exit could correlate with both technology investment and forecast performance. To confirm that compositional churn does not drive the baseline, we re-estimate the baseline on the 66 brokerages that issue at least one forecast in every fiscal year from 2012 to 2022 (798,281 forecasts). Table OA4 mirrors the six-column design of Table 5. The *AI Adoption* coefficient is nearly identical to the full-sample estimate in every column (0.027 with analyst-firm fixed effects and controls, against 0.025 in the full sample), so the baseline association is not an artifact of brokerages entering or exiting the sample.

Excluding Likely Off-the-Shelf AI Adopters

Our hiring-based measure does not observe AI adopted through outsourcing or off-the-shelf services. If such adopters populate the comparison group, they attenuate the estimated association, which is one reason we interpret the baseline as a lower bound (Section III). Here we take that argument to the data by excluding from the comparison group the brokerages most likely to be unmeasured adopters, defined as those that post no AI-related research positions over the 2010–2022 posting window while their non-AI technology hiring share lies in the sample's top quartile, the profile most consistent with a technologically active brokerage buying AI externally. Fifteen brokerages, accounting for 41,139 forecasts (about 5 percent of the sample), are excluded. Table OA5 shows that the estimates are unchanged or slightly larger in every column (0.026, against 0.025 in the full sample, with analyst-firm fixed effects and controls), consistent with the lower-bound interpretation. Stricter definitions, requiring at least ten research-related postings over the window or classifying by posting counts, deliver the same inference (untabulated; the corresponding coefficients range from 0.025 to 0.026 with t-statistics of 2.95 to 3.10).

Controlling for Big-Data Hiring

AI hiring could in principle proxy for a brokerage's broader investment in data infrastructure rather than for AI specifically; relatedly, Chi et al. (2025) document that brokerages' big-data adoption improves their analysts' forecasts. To separate the two, we construct *BigData Adoption* exactly parallel to *AI Adoption*, defined as the trailing four-quarter average of the brokerage's share of research-related postings that require a big-data skill but no AI skill (for example, Big Data, Apache Hadoop, Apache Spark, Apache Kafka, NoSQL databases, and Teradata), min-max standardized within firm-year. Table OA6 adds this control to the baseline ladder. The *AI Adoption* coefficient is essentially unchanged in every column (0.024, $t = 2.75$, with analyst-firm fixed effects and controls), while *BigData Adoption* is significant only in the columns without controls and statistically indistinguishable from zero once they enter (0.005, $t = 0.56$). The accuracy association loads on AI specifically rather than on generic big-data capability.

Brokerages Covered in Shanthikumar and Yoo (2026) versus the Remaining Brokerages

Concurrent work by Shanthikumar and Yoo (2026) studies sixteen large brokerages. To assess whether our estimates are specific to that segment of the industry, Table OA7 re-estimates the baseline separately on the brokerages they cover (Panel A) and on the remaining 63 brokerages (Panel B); the two groups split the forecast sample roughly in half. The association is present in both. With analyst-firm fixed effects and controls, the coefficient is 0.028 ($t = 2.35$) among the covered brokerages and 0.024 ($t = 1.87$) among the remaining ones. Our results are neither confined to, nor driven by, the largest brokerages.

OA III. ALTERNATIVE ERROR METRIC AND LAG STRUCTURE

Decision Fatigue with Absolute Forecast Errors

Table OA3, Panel C, establishes that the baseline relation holds with the absolute forecast error (*AFE*) as the dependent variable under firm-year fixed effects. Table OA8 applies the same design to the decision-fatigue analysis—the dependent variable is *AFE*, every specification includes firm-year fixed effects, and analyst and analyst-firm fixed effects are added across columns. The interaction between *AI Adoption* and *Fcast Under Fatigue* is negative in all six columns and statistically significant in five of them, indicating that brokerage AI reduces forecast errors by more precisely when the analyst is fatigued, matching the relative-accuracy results of Table 8.

The Lag Structure of the Baseline Relation

Section V reports that the accuracy benefit tracks a brokerage's recent AI capability. Table OA9 presents the underlying estimates. Replacing *AI Adoption* with its value lagged one to four quarters, the association is positive and significant at the one- and two-quarter lags (0.027 and 0.022) and statistically indistinguishable from zero at three and four quarters. When the current four-quarter window and the non-overlapping window ending four quarters earlier enter the regression jointly, the current window retains the baseline magnitude (0.025, t = 2.77) while the year-old window contributes nothing (-0.001, t = -0.15), consistent with AI capability that depreciates unless renewed by continued hiring.

REFERENCES

Chi, F., B.-H. Hwang, and Y. Zheng. 2025. The use and usefulness of big data in finance: Evidence from financial analysts. *Management Science* 71 (6): 4599–4621.

Clement, M. B. 1999. Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics* 27 (3): 285–303. https://doi.org/10.1016/s0165-4101(99)00013-0.

Jacob, J., T. Z. Lys, and M. A. Neale. 1999. Expertise in forecasting performance of security analysts. *Journal of Accounting and Economics* 28 (1): 51–82. https://doi.org/10.1016/s0165-4101(99)00016-6.

Shanthikumar, D. M. and I. S. Yoo. 2026. Beyond automation: AI and the human value of sell-side analysts (Working paper). https://ssrn.com/abstract=5040339.

FIGURE OA1. Validating the AI Adoption Measure

Panel A: AI Patenting

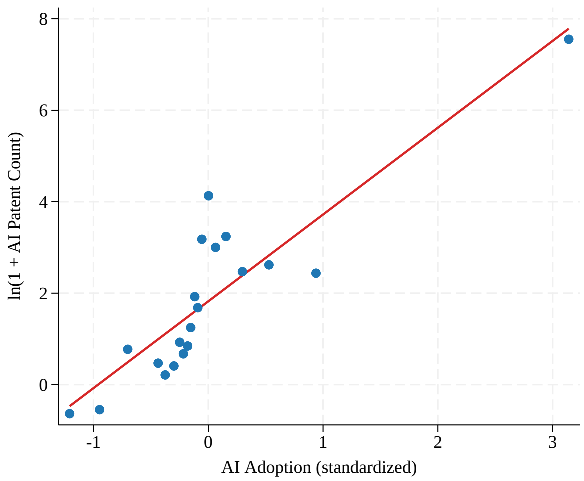

Panel B: AI Disclosure

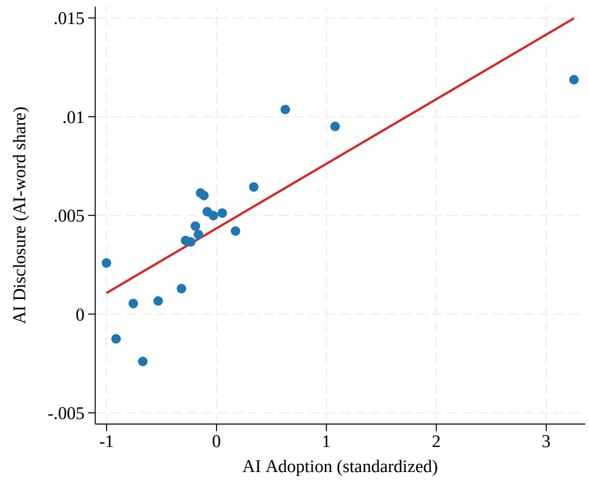

This figure visualizes the validation reported in Table 3. Each panel is a binned scatter plot of an independent measure of brokerage AI activity against the hiring-based *AI Adoption* measure (standardized), after partialing out year-quarter fixed effects: broker-quarter observations are sorted into 20 equal-count bins and the line is a linear fit. The raw shares and counts are non-negative; the slightly negative binned values arise only from demeaning each axis by its year-quarter average. **Panel A** plots AI patenting ($\ln(1+\mathit{AI\ Patent\ Count})$) for the 16 brokerages matched to USPTO filings; **Panel B** plots AI disclosure (*AI Disclosure*) for the 24 publicly listed brokerages. In both panels the relationship is positive: brokerages that hire more AI also patent and disclose more AI, corroborating the hiring-based measure.

FIGURE OA2. AI Adoption and Relative Forecast Accuracy

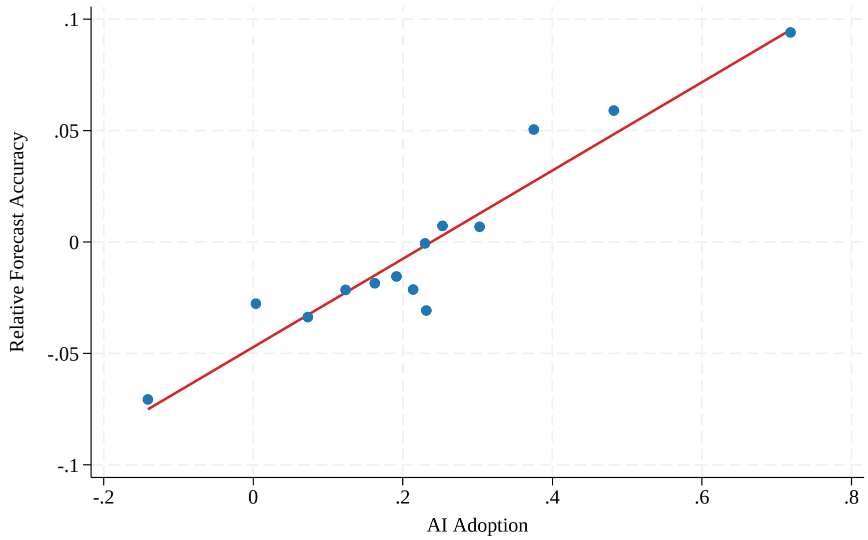

This figure is a binned scatter plot of analysts' relative forecast accuracy (*Relative Accuracy*) against their brokerage's AI adoption (*AI Adoption*). Both variables are taken net of analyst-firm means (the within transformation underlying the analyst-firm fixed-effects specification in Table 5, Column 5), so the slope of the fitted line corresponds to the within (analyst-firm) estimate; *Relative Accuracy* is centered at zero and forecasts are sorted into 20 equal-count bins by *AI Adoption*. The positive slope shows that, comparing forecasts by the same analyst on the same firm over time, those issued when the brokerage has adopted more AI are more accurate relative to peers covering the same firm. Table 5 reports the corresponding regressions, including the specification with the full set of controls.

TABLE OA1. I/B/E/S Brokerages with the Highest AI Adoption

Rank	Brokerage	I/B/E/S Estimator Code (ESTIMID)	*AI Adoption* (%)	Headquarters
1	Goldman Sachs	GOLDMAN	4.0	New York, NY
2	Barclays Capital	FRCLAYSC	3.7	New York, NY
3	J.P. Morgan	JPMORGAN	3.4	New York, NY
4	Credit Suisse	FBOSTON	3.1	New York, NY
5	William Blair & Company	BLAIR	2.9	Chicago, IL
6	Morgan Stanley	MORGAN	2.7	New York, NY
7	Bank of America Securities	MERRILL	2.7	New York, NY
8	Wells Fargo Securities	WHEAT	2.4	Charlotte, NC
9	BMO Capital Markets	BURNS	2.2	New York, NY
10	Jefferies	JEFFEREG	2.0	New York, NY

This table lists the ten most AI-adopting brokerages in the sample. *AI Adoption* is the brokerage's trailing-four-quarter AI-related hiring share, in percent (the share of its job postings classified as AI-related), averaged over its quarters in the 2012–2022 sample. *I/B/E/S code* is the brokerage's I/B/E/S identifier. The ranking is restricted to brokerages with more than ten sample analysts, which removes a small energy-research boutique whose AI share is driven by its parent organization's hiring rather than its research unit. *Headquarters* is the brokerage's primary U.S. research office. All ten are full-service brokerages with broad industry coverage.

TABLE OA2. Additional Analyses

Panel A: Analyst Experience Interactions

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.011	0.006	0.013	0.010	0.007	0.009
	(1.24)	(0.65)	(1.26)	(1.13)	(0.80)	(0.89)
AI Adoption × *Low Firm Experience*	0.025**	0.027***	0.035***			
	(2.46)	(2.87)	(2.77)			
Low Firm Experience	-0.024***	-0.014***	-0.015**			
	(-5.26)	(-3.59)	(-2.46)			
AI Adoption × *Low General Experience*				0.028***	0.030***	0.041***
				(2.75)	(3.09)	(3.20)
Low General Experience				-0.000	0.002	0.003
				(-0.07)	(0.31)	(0.36)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	816,418	816,244	813,205	816,418	816,244	813,205
Adj. R^2	0.207	0.223	0.269	0.207	0.223	0.269
Fixed Effects	No	Analyst	Analyst×firm	No	Analyst	Analyst×firm

Panel B: Forecast Characteristics

Dependent variable:	*Frequency*		*Thoroughness*		*Timeliness* (0/1)	
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.023***	0.024***	0.025***	0.027***	0.101***	0.059***
	(4.02)	(4.60)	(3.47)	(3.89)	(4.61)	(2.60)
Model:	OLS	OLS	OLS	OLS	Logit	Logit
Controls	No	Yes	No	Yes	No	Yes
Observations	492,885	486,055	492,885	486,055	333,786	330,960
Adj./Pseudo R^2	0.304	0.327	0.461	0.465	0.001	0.046
Fixed Effects	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm	Analyst×firm

Panel C: Forecast Optimism

Dependent variable:	*Relative Optimism*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	-0.003	-0.012	0.034***	0.014*	0.005	0.004
	(-0.29)	(-0.74)	(4.97)	(1.92)	(0.65)	(0.52)
Controls	No	Yes	No	Yes	No	Yes
Observations	814,464	814,464	814,291	814,291	811,237	811,237
Adj. R^2	0.000	0.002	0.082	0.086	0.236	0.238
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm

Panel D: Alternative Fixed Effects

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.141***	0.037***	0.193***	0.018**	0.526***	0.016
	(15.98)	(4.54)	(17.03)	(2.36)	(17.53)	(1.17)
Controls	No	Yes	No	Yes	No	Yes
Observations	816,418	816,418	814,567	816,244	792,042	792,042
Adj. R^2	0.003	0.209	0.042	0.223	0.104	0.349
Fixed Effects	Broker	Broker	Broker×firm	Broker×firm	Analyst×firm×year	Analyst×firm×year

Panel E: Accuracy at the Brokerage-Quarter Level

Dependent variable:	*Average Relative AFE*			
	(1)	(2)	(3)	(4)
AI Adoption	-0.025***	-0.019**	-0.028***	-0.021*
	(-3.70)	(-1.99)	(-4.02)	(-1.95)
Controls	No	Yes	No	Yes
Quarter FE	No	No	Yes	Yes
Observations	2,916	2,916	2,916	2,916
Within R^2	0.001	0.013	0.001	0.014

This table presents five sets of additional analyses. The control variables are those in Table 5. **Panel A** interacts *AI Adoption* with indicators for below-median firm-specific experience (*Low Firm Experience*, Columns 1–3) and general experience (*Low General Experience*, Columns 4–6) in the baseline specification (Equation 5); each indicator equals one if the analyst's experience is at or below the median among analysts covering the same firm-year. **Panel B** relates *AI Adoption* to *Frequency* and *Thoroughness* (OLS, Columns 1–4) and *Timeliness* (conditional logit, Columns 5–6): *Frequency* is the number of annual EPS forecasts the analyst issues for the firm in the quarter and *Thoroughness* the number of distinct forecast horizons, each entered as a within-firm-quarter rank scaled between zero and one; *Timeliness* equals one if the analyst's first revision after the firm's earnings announcement is issued within two days. **Panel C** re-estimates the baseline specification with *Relative Optimism* (Equation OA.1) as the dependent variable. **Panel D** re-estimates Equation (5) using alternative fixed-effects structures. **Panel E** conducts the accuracy test at the brokerage-quarter level, on the constant sample of 2,916 brokerage-quarters with nonmissing controls, identical to the estimation sample of Table 4. *Average Relative AFE* is the brokerage-quarter mean of the relative *AFE*: each forecast's absolute error, scaled by beginning-of-year price, minus the mean price-scaled error across all analysts' forecasts for the same firm in the same quarter, averaged across the brokerage's forecasts in the quarter. The within-firm-quarter adjustment removes variation arising from differences in coverage portfolios, and negative coefficients indicate errors below the consensus error on the same firms in the same periods. The controls are the brokerage-level determinants of Table 4 other than *Peer Adoption*, with the regressors non-scaled (i.e., not min-max standardized). The sample includes all annual analyst EPS forecasts over 2012–2022. The remaining variables are defined in Appendix A. t-statistics (z-statistics in Columns 5–6 of Panel B) based on standard errors clustered by analyst (by brokerage in Panel E) are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE OA3. Using AFE as Dependent Variable: Baseline Results

Dependent variable:	*AFE*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	-0.032***	-0.023**	-0.046***	-0.020**	-0.052***	-0.020**
	(-3.22)	(-2.06)	(-4.80)	(-2.15)	(-4.65)	(-1.96)
Controls	No	Yes	No	Yes	No	Yes
Firm×Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Analyst FE	No	No	Yes	Yes	No	No
Analyst×Firm FE	No	No	No	No	Yes	Yes
Observations	775,129	775,129	774,947	774,947	771,357	771,357
Adj. R^2	0.840	0.842	0.841	0.843	0.848	0.850

This table re-estimates the baseline relation of Equation (5) using the absolute forecast error (*AFE*) as the dependent variable: the absolute difference between the forecast and realized EPS, scaled by the stock price at the start of the fiscal year and winsorized at the 1st and 99th percentiles, with the regressors non-scaled (i.e., not min-max standardized). The sample includes all annual analyst EPS forecasts over 2012–2022. The control variables are those in Table 5. Because *Relative Accuracy* is normalized within firm-year, every column absorbs that component through firm-year fixed effects, adding analyst fixed effects in Columns 3–4 and analyst-firm fixed effects in Columns 5–6; a negative coefficient indicates that AI adoption reduces forecast error, and *Analyst Coverage* is subsumed by the firm-year fixed effects. The remaining variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE OA4. Continuously-Active Brokerages

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.069***	0.024***	0.142***	0.019**	0.195***	0.027***
	(10.24)	(3.24)	(15.25)	(2.48)	(16.29)	(3.13)
<u>*Control variables*</u>						
Firm Experience		0.038***		0.021***		0.046***
		(5.74)		(3.73)		(3.65)
General Experience		0.000		0.064***		0.066***
		(0.05)		(4.81)		(3.35)
Effort		-0.057***		-0.067***		-0.075***
		(-9.42)		(-12.44)		(-12.05)
Broker Size		-0.023***		-0.008		0.029**
		(-2.74)		(-0.87)		(2.08)
Forecast Age		-1.401***		-1.411***		-1.431***
		(-174.28)		(-176.81)		(-174.35)
Firms Followed		0.020***		0.002		0.003
		(2.64)		(0.27)		(0.31)
Analyst Coverage		0.004		0.005		0.032***
		(1.08)		(1.36)		(3.72)
Observations	798,281	798,281	798,108	798,108	795,167	795,167
Adj. R^2	0.001	0.207	0.015	0.222	0.050	0.268
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm

This table re-estimates the baseline relation of Table 5 on the subsample of continuously active brokerages, the 66 of the 79 brokerages in the sample that issue at least one forecast in every fiscal year from 2012 to 2022 (798,281 forecasts). All variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE OA5. Excluding Likely Off-the-Shelf AI Adopters from the Comparison Group

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.071***	0.024***	0.142***	0.019**	0.194***	0.026***
	(10.36)	(3.28)	(15.18)	(2.45)	(16.17)	(3.10)
<u>*Control variables*</u>						
Firm Experience		0.036***		0.017***		0.040***
		(5.27)		(3.04)		(3.16)
General Experience		0.002		0.068***		0.073***
		(0.21)		(4.92)		(3.68)
Effort		-0.058***		-0.068***		-0.075***
		(-9.48)		(-12.40)		(-11.87)
Broker Size		-0.022**		-0.009		0.028**
		(-2.53)		(-0.96)		(1.98)
Forecast Age		-1.404***		-1.413***		-1.433***
		(-171.66)		(-173.78)		(-171.23)
Firms Followed		0.018**		0.002		0.003
		(2.33)		(0.22)		(0.29)
Analyst Coverage		0.005		0.005		0.033***
		(1.17)		(1.27)		(3.80)
Observations	775,279	775,279	775,113	775,113	772,319	772,319
Adj. R^2	0.001	0.208	0.015	0.222	0.051	0.268
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm

This table re-estimates the baseline relation of Table 5 after excluding from the comparison group the brokerages most likely to have adopted AI through channels invisible to the hiring-based measure. A brokerage is excluded if it posts no AI-related research positions over the 2010–2022 posting window while its non-AI technology hiring share (non-AI technology postings as a share of its research-related postings, pooled over the window) falls in the top quartile of the sample. Fifteen brokerages, accounting for 41,139 forecasts (about 5 percent of the sample), are excluded. All variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE OA6. Controlling for Brokerages' Big-Data Hiring

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.048***	0.022***	0.116***	0.017**	0.169***	0.024***
	(6.15)	(2.77)	(12.17)	(2.19)	(14.21)	(2.75)
BigData Adoption	0.033***	0.005	0.090***	0.003	0.131***	0.005
	(4.30)	(0.68)	(9.95)	(0.38)	(11.99)	(0.56)
<u>*Control variables*</u>						
Firm Experience		0.038***		0.022***		0.047***
		(5.76)		(3.93)		(3.76)
General Experience		0.002		0.068***		0.071***
		(0.17)		(4.99)		(3.62)
Effort		-0.054***		-0.066***		-0.073***
		(-9.09)		(-12.32)		(-11.78)
Broker Size		-0.024***		-0.010		0.021
		(-2.78)		(-1.15)		(1.49)
Forecast Age		-1.404***		-1.414***		-1.434***
		(-175.94)		(-178.48)		(-175.85)
Firms Followed		0.018**		-0.000		-0.000
		(2.31)		(-0.04)		(-0.03)
Analyst Coverage		0.005		0.005		0.032***
		(1.31)		(1.43)		(3.76)
Observations	811,416	811,416	811,244	811,244	808,239	808,239
Adj. R^2	0.001	0.208	0.015	0.223	0.051	0.269
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm

This table re-estimates the baseline relation of Table 5 adding the brokerage's big-data hiring as a control. *BigData Adoption* is constructed exactly parallel to *AI Adoption*: the trailing four-quarter average of the brokerage's share of research-related postings that require a big-data skill but no AI skill, min-max standardized within firm-year. Big-data skills comprise the distributed data-infrastructure stack (for example, Big Data, Apache Hadoop, Apache Spark, Apache Kafka, NoSQL databases, and Teradata), identified from the posted skill requirements in the same Lightcast data. All other variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE OA7. Brokerages Covered in Shanthikumar and Yoo (2026) versus the Remaining Brokerages

Panel A: Brokerages Covered in Shanthikumar and Yoo (2026)

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.076***	0.011	0.149***	0.016	0.213***	0.028**
	(8.29)	(1.07)	(11.91)	(1.54)	(13.54)	(2.35)
Controls	No	Yes	No	Yes	No	Yes
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm
Observations	400,601	400,601	400,505	400,505	399,100	399,100
Adj. R^2	0.001	0.201	0.014	0.215	0.052	0.263

Panel B: The Remaining Brokerages

Dependent variable:	*Relative Accuracy*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	0.094***	0.047***	0.140***	0.025**	0.173***	0.024*
	(8.80)	(4.43)	(9.60)	(2.08)	(9.44)	(1.87)
Controls	No	Yes	No	Yes	No	Yes
Fixed Effects	No	No	Analyst	Analyst	Analyst×firm	Analyst×firm
Observations	415,817	415,817	415,734	415,734	413,971	413,971
Adj. R^2	0.001	0.214	0.016	0.231	0.050	0.276

This table re-estimates the baseline relation of Table 5 on two subsamples of brokerages. **Panel A** restricts the sample to forecasts issued by analysts at the 16 brokerages examined in Shanthikumar and Yoo (2026); **Panel B** uses the remaining 63 brokerages. Each column matches the corresponding specification of Table 5, with the controls of that table added in the even-numbered columns. All variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE OA8. Using AFE as Dependent Variable: Decision Fatigue

Dependent variable:	*AFE*					
	(1)	(2)	(3)	(4)	(5)	(6)
AI Adoption	-0.027***	-0.018*	-0.043***	-0.016*	-0.050***	-0.018*
	(-3.01)	(-1.71)	(-4.47)	(-1.75)	(-4.33)	(-1.65)
AI Adoption × *Fcast Under Fatigue*	-0.033**	-0.037***	-0.023**	-0.029***	-0.013	-0.020*
	(-2.57)	(-3.04)	(-2.12)	(-2.73)	(-1.24)	(-1.83)
Fcast Under Fatigue	0.001	0.002***	0.000	0.001	-0.001	0.000
	(0.86)	(2.83)	(-0.33)	(1.50)	(-1.17)	(0.57)
Controls	No	Yes	No	Yes	No	Yes
Firm×Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Analyst FE	No	No	Yes	Yes	No	No
Analyst×Firm FE	No	No	No	No	Yes	Yes
Observations	775,129	775,129	774,947	774,947	771,357	771,357
Adj. R^2	0.840	0.842	0.841	0.843	0.847	0.850

This table re-estimates the decision-fatigue analysis of Table 8 using the absolute forecast error (*AFE*, as defined in Table OA3) as the dependent variable, under the same fixed-effects specifications as that table. *Fcast Under Fatigue* equals one if the forecast is the fourth or later forecast issued by the same analyst on the same day and is issued between 9:00 and 23:00 (the ordering is computed over the full sample). A negative interaction indicates that AI adoption reduces forecast error by more when the analyst is fatigued. Odd-numbered columns exclude controls; even-numbered columns add the controls of Table OA3, non-scaled (i.e., not min-max standardized). The remaining variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

TABLE OA9. Persistence of the AI Adoption–Accuracy Relationship: Lagged AI Adoption

Dependent variable:	*Relative Accuracy*				
	(1)	(2)	(3)	(4)	(5)
$AI\ Adoption_t$					0.025***
					(2.77)
$AI\ Adoption_{t\text{-}1}$	0.027***				
	(2.98)				
$AI\ Adoption_{t\text{-}2}$		0.022***			
		(2.58)			
$AI\ Adoption_{t\text{-}3}$			0.006		
			(0.72)		
$AI\ Adoption_{t\text{-}4}$				0.003	-0.001
				(0.28)	(-0.15)
Controls	Yes	Yes	Yes	Yes	Yes
Analyst×Firm FE	Yes	Yes	Yes	Yes	Yes
Observations	813,205	810,474	808,618	807,502	806,482
Adj. R^2	0.269	0.269	0.269	0.269	0.269

This table examines whether the relationship between brokerage AI adoption and forecast accuracy persists over time. The dependent variable is *Relative Accuracy*, and the regressors are constructed as

$$AI\ Adoption_{t-k} = \frac{1}{4}\sum_{s=0}^{3} AI\ Hiring_{t-k-s},$$

where *AI Hiring* is the brokerage's quarterly share of research postings requiring an AI skill, so that $AI\ Adoption_{t-k}$ is its four-quarter trailing average ending k quarters before the forecast quarter, min-max standardized within firm-year. All columns include the same controls and analyst-firm fixed effects as the baseline forecast-accuracy specification. Columns 1–4 enter each lag separately. Column 5 enters $AI\ Adoption_t$ and $AI\ Adoption_{t\text{-}4}$ jointly; their four-quarter windows do not overlap, so the two are not mechanically collinear. The remaining variables are defined in Appendix A. t-statistics based on standard errors clustered by analyst are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.