Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI Washing and the Informational Content of Analyst Consensus

Yingying Yang
University of Southampton
Nan Li
University of Southampton
Yiqi Wang
University of Southampton
Ziting Wang
Ziting.Wang@soton.ac.uk
University of Southampton

Research Article

Keywords: AI washing, corporate disclosure, analysts forecast disagreement, false consensus

Posted Date: November 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11614769/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

AI Washing and the Informational Content of Analyst Consensus

ABSTRACT

We investigate whether corporate AI narratives that are weakly substantiated by subsequent patent-based technological output are associated with less informative analyst consensus. Using 13,450 U.S. firm-years from 2014 to 2023, we measure ex post AI narrative–output decoupling as the difference between within-year percentile ranks of AI-related language in 10-K filings and AI-classified pre-grant patent publications observed after the initial post-filing analyst consensus. Higher decoupling is associated with lower analyst forecast dispersion but greater consensus forecast error. A one-standard-deviation increase corresponds to differences equivalent to 2.97% and 3.10% of the respective sample means. Event-time analyses indicate that the dispersion relationship emerges after the filing. Analyst-level evidence shows greater convergence toward prevailing forecasts, more herding, less bold forecasting, and lower model-implied private-information precision, with weaker associations among more experienced analysts. By contrast, high AI narrative intensity accompanied by positive patent-based AI output is associated with both lower disagreement and lower forecast error. The findings suggest that apparently stronger analyst consensus may not signal a better information environment when corporate technological claims are difficult to verify.

JEL codes: M41; G14; G24; D82

Keywords: AI washing; corporate disclosure; analysts forecast disagreement; false consensus

Highlights

- AI washing is associated with lower analyst disagreement but higher forecast error.
- Narrative–substance decoupling is associated with less informative consensus.
- Analysts exhibit greater herding and lower private-information precision.
- The associations are weaker among more experienced analysts.
- High AI narrative intensity accompanied by positive patent-based AI output is associated with both lower forecast disagreement and lower forecast error.

1 Introduction

Financial analysts play a central role in converting corporate disclosure into market expectations. A large literature shows that more informative and easier-to-process disclosure is associated with lower forecast dispersion and greater forecast accuracy (Lang and Lundholm 1996; Healy and Palepu 2001; Beyer et al. 2010; Lehavy et al. 2011). Recent evidence also cautions against equating greater corporate disclosure with higher-quality information. Vo and Pham (2025), for example, show that firms facing greater cybersecurity risk issue more management forecasts in response to heightened information demand, yet those forecasts are more optimistic and exhibit lower precision and accuracy. Yet agreement is not itself evidence of information quality. Analyst forecasts combine common and analyst-specific information, and greater reliance on common information can reduce dispersion without increasing the precision of the beliefs analysts share (Barron et al. 1998). Prior work further shows that analysts may herd or strategically compress disagreement (Clement and Tse 2005; Liu and Natarajan 2012). Much less is known about a different source of uninformative consensus: whether the corporate disclosure being processed can itself induce analysts to agree more without becoming better informed. We examine this possibility using corporate disclosures about artificial intelligence (AI).

AI is a particularly useful setting because the underlying activity is economically important yet difficult to verify contemporaneously. Bhimani et al. (2026) argue that digitalisation is reshaping the structure and sources of accounting-relevant information and challenging conventional assumptions about the data used for organisational decision-making and control. Firms' AI capabilities can reside in specialized human capital, proprietary data, software, models, and complementary organizational investments that are only imperfectly observable to outsiders (Babina et al. 2024; Barrios et al. 2026). AI disclosure may therefore be informative on average while differing substantially in the extent to which the narrative is subsequently supported by observable technological activity. This distinction parallels the corporate-washing literature, which focuses on divergence between external claims and underlying

outcomes rather than disclosure intensity alone (Kim and Lyon 2015; Baker et al. 2024). We refer to such divergence as AI washing, or AI narrative–output decoupling. The construct does not presume intentional deception. Our question is narrower: when analysts observe a common AI narrative but cannot yet verify its technological substance, do disclosures that are subsequently less substantiated generate greater agreement without greater forecast quality?

We study 13,450 U.S. firm-year observations from 2014–2023. We measure AI narrative intensity from the MD&A section of the original Form 10-K and subsequently verifiable technological output using AI-classified pre-grant patent publications from the USPTO Artificial Intelligence Patent Dataset. Following the disclosure–substance approach of Baker et al. (2024), we rank the narrative and output components separately within year and define AIWashing as the difference between the two percentile ranks. The timing is central. The AI narrative is publicly available before analysts form the first post-10-K FY1 consensus, whereas the patent output used to evaluate subsequent technological substantiation is measured only after that consensus. AIWashing is therefore an ex post classification of the focal disclosure, not information available to analysts when they form the forecasts we examine.

We document a sharp divergence between agreement and accuracy. A one-standard-deviation increase in AIWashing is associated with a 2.97% reduction in forecast disagreement relative to its sample mean, but a 3.10% increase in consensus forecast error. The timing of the convergence is informative: AIWashing is not meaningfully related to disagreement before the focal 10-K filing, whereas the association becomes increasingly negative after the filing. Moreover, the initially tighter consensus is followed by larger forecast revisions and a greater likelihood that subsequent revisions move toward realized earnings. We refer to this combination—greater contemporaneous agreement without greater accuracy—as false consensus. Complementary instrumental-variable evidence based on firms' historical exposure to AI-strong university talent yields the same directional inference, although we interpret it subject to the maintained exclusion restriction.

The analyst-level evidence is consistent with differences in the information composition underlying forecast consensus. Greater AIWashing is associated with stronger movement toward the prevailing leave-one-out consensus, more herding, less bold forecasting, and lower model-implied private-information precision. These associations are weaker for analysts with greater firm-specific experience and, in supplementary analyses, for analysts with stronger prior forecast performance, larger brokerage resources, and greater industry specialization. This pattern is consistent with analysts relying more heavily on a shared corporate narrative when their capacity to independently evaluate it is weaker. Importantly, the evidence differs for substantive AI activity. Strong AI narratives accompanied by observable AI output are associated with both lower disagreement and lower forecast error, whereas narrative–output decoupling is associated with lower disagreement but higher error. Greater AIWashing is also associated with larger subsequent earnings-announcement reactions, indicating that the apparent resolution of disagreement does not fully resolve uncertainty about fundamentals.

Our study makes two contributions. First, we extend the literature on corporate disclosure and analyst information environments by documenting a divergence between forecast agreement and accuracy. Prior research establishes that disclosure shapes analysts' forecast properties and their common and private information (Lang and Lundholm 1996; Barron et al. 1998, 2002, 2017; Chen et al. 2010; Lehavy et al. 2011), while the corporate-washing literature examines discrepancies between external claims and underlying outcomes (Kim and Lyon 2015; Baker et al. 2024; Barrios et al. 2026). We connect these perspectives by showing that AI narratives weakly substantiated by subsequent patent-based output are associated with lower analyst disagreement but higher forecast error. Second, extending research on analyst herding and information production (Clement 1999; Clement and Tse 2005; Liu and Natarajan 2012; Brown et al. 2015), we document stronger convergence toward prevailing forecasts, greater herding, and lower model-implied private-information precision, with weaker associations among more experienced analysts (Mikhail et al. 1997). These findings

qualify the interpretation of analyst consensus as an indicator of information quality when corporate technological claims are difficult to verify.

2 Analysts, Disclosure, and the Information Environment

2.1 Corporate AI Disclosure and Narrative–Output Decoupling

Corporate disclosure is particularly important when economically relevant activities are not fully reflected in financial statements or are costly for outsiders to verify (Healy and Palepu 2001; Beyer et al. 2010). These frictions are pronounced for innovation and intangible investment because expenditures on knowledge, human capital, and internally generated technology are only imperfectly captured by conventional accounting measures (Lev and Sougiannis 1996; Peters and Taylor 2017). These challenges are becoming more pronounced as digitalisation changes the nature, structure, and sources of economically relevant organisational data. Bhimani et al. (2026) argue that digital technologies are reshaping accounting information, decision processes, and control arrangements, requiring researchers to look beyond conventional transaction-based financial information. Prior research consequently finds greater uncertainty surrounding firms with substantial technology-related intangibles and an important role for analysts in supplementing the financial reporting system (Barron et al. 2002a). AI magnifies these informational frictions because economically meaningful AI capabilities can reside in specialized labor, data, software, models, and complementary organizational investments that are difficult to observe separately (Babina et al. 2024; Barrios et al. 2026). Corporate AI narratives can therefore convey useful information about otherwise opaque technological activity, but the underlying substance of that activity is difficult to verify contemporaneously. Barrios et al. (2026), for example, document that AI disclosure is informative about subsequent AI investment and operating outcomes on average, while also emphasizing the intangible and opaque nature of the underlying investment.

Informative disclosure on average does not imply that disclosure intensity is uniformly aligned with underlying substance. Research on corporate washing instead

emphasizes the correspondence between external claims and observable underlying outcomes. Kim and Lyon (2015) distinguish firms' environmental communication from their environmental performance, while Baker et al. (2024) identify diversity washing from the discrepancy between firms' diversity-related disclosure and workforce diversity. Barrios et al. (2026) similarly distinguish firms whose AI disclosures are supported by AI-related human-capital investment from firms for which such support is weak. Building on this logic, we define AI washing as narrative–output decoupling rather than as AI disclosure itself or simply low AI activity: a larger gap indicates that the relative prominence of AI in corporate communication exceeds subsequently verifiable technological output. This definition does not require intentional deception; the gap may also reflect optimism, implementation delays, failed projects, or other forms of ex post non-substantiation.

2.2 Public Disclosure, Common Information, and Analyst Beliefs

Financial analysts jointly process corporate disclosures and information generated through their own research (Beyer et al. 2010; Brown et al. 2015). A central finding of the disclosure literature is that more informative and easier-to-process disclosure improves analysts' information environment. Firms with more informative disclosure policies have more accurate and less dispersed forecasts (Lang and Lundholm 1996), while the consequences of 10-K readability and complexity for the information environment can vary with the evaluation horizon (Lehavy et al. 2011; Davis et al. 2026). More recently, Park and Zach (2025) show that unmet analyst information demand generates persistent uncertainty, and Chircop et al. (2026) document that improved dissemination of economically relevant information increases analyst forecast accuracy. Recent evidence further shows that greater discretion in the presentation of accounting information can impair analysts' information production. Ha et al. (2025) find that the removal of a standardized operating-income definition under IFRS reduced analysts' propensity to issue operating-income forecasts and lowered forecast accuracy, effects that dissipated after regulatory intervention restored reporting standardization. These studies establish an important benchmark: the analyst

information environment should be evaluated using both forecast disagreement and forecast accuracy (Lehavy et al. 2011; Lu et al. 2026).

That benchmark does not imply that public disclosure mechanically produces better consensus. Analysts' forecasts combine common and analyst-specific information, and these components have distinct implications for disagreement and accuracy (Barron et al. 1998). Accounting disclosures can also alter those components differently. Barron, Byard, and Kim (2002b) find that earnings announcements can stimulate idiosyncratic information production, whereas Barron, Byard, and Yu (2017) show that balance-sheet and segment disclosures increase private-information precision while management forecasts primarily increase common information. Chen, Cheng, and Lo (2010) likewise distinguish analysts' information-discovery role from their interpretation of public disclosure and show that interpretation becomes more important when accounting information is difficult to process. More generally, public information can interact with heterogeneous prior information and differential interpretation rather than simply eliminate disagreement (Kim and Verrecchia 1991, 1994). Accordingly, the effect of disclosure on consensus depends not only on how much information is disclosed, but on its precision, verifiability, and the extent to which it induces common versus analyst-specific information processing.

2.3 AI Washing and Analyst Forecast Convergence

The effect of opaque technological disclosure on analyst disagreement is therefore theoretically ambiguous. Opacity by itself can increase analyst-specific information production: Barron et al. (2002a) find that analysts place greater relative weight on private information for firms with substantial high-technology intangibles, resulting in lower consensus. Difficult corporate information can likewise increase analysts' interpretation role (Chen et al. 2010), and public disclosure can generate heterogeneous informed judgments rather than uniform updating (Kim and Verrecchia 1994). In contrast, managerial disclosures can increase analysts' common information when they provide a shared signal about future performance (Barron et al. 1998; Barron et al.

2017). The distinction is important. Technological opacity alone need not generate convergence; convergence becomes more plausible when opacity is accompanied by a salient common narrative that analysts can observe but cannot readily validate against contemporaneous hard evidence.

AI narrative–output decoupling creates precisely this configuration. Analysts observe the same AI-oriented disclosure, while the technological substance required to evaluate that narrative is only imperfectly observable at the time forecasts are formed. When independently verifiable evidence is limited, the common narrative can receive greater relative weight in posterior beliefs, causing analysts' forecasts to co-move even if the signal itself is not highly precise (Barron et al. 1998; Barron et al. 2017). Importantly, this argument does not require analysts to identify AI washing contemporaneously: narrative–output decoupling is an ex post characterization of disclosures that differ in their subsequent technological substantiation. We therefore predict that disclosures subsequently characterized by greater narrative–output decoupling are associated with stronger contemporaneous forecast convergence.

H1: *Greater ex post AI narrative–output decoupling is associated with lower analyst forecast disagreement following the focal corporate disclosure.*

2.4 False Consensus: Convergence without Improved Forecast Quality

Forecast agreement and forecast quality are not equivalent. In the Barron et al. (1998) framework, dispersion reflects differences across analysts' beliefs but does not capture errors contained in the information they share. Analysts can therefore exhibit low dispersion while collectively placing weight on an imprecise common signal. This distinction is reinforced by evidence on analysts' strategic forecasting behavior. Liu and Natarajan (2012) show that observed forecast dispersion can substantially understate the dispersion of unmanaged forecasts, while Clement and Tse (2005) find that herding forecasts are less accurate and incorporate private information less completely than bold forecasts. Related evidence from the strategic-disclosure literature shows that disclosure quantity and disclosure quality can move in opposite directions. Vo and

Pham (2025) document that firms facing heightened cybersecurity risk issue more management forecasts even though those forecasts are less precise and less accurate. What remains unknown is whether such informationally imperfect corporate communication also changes the cross-sectional structure of analysts' beliefs, causing analysts to agree more rather than less. Thus, declining disagreement is consistent with an improved information environment only when the information generating that agreement is itself sufficiently precise (Lang and Lundholm 1996; Lehavy et al. 2011; Barron et al. 1998). More broadly, recent evidence indicates that analyst forecast quality deteriorates when the underlying reporting environment provides less standardized or less decision-useful information (Ha et al., 2025).

This distinction creates the central theoretical gap addressed here. The disclosure literature largely associates more informative disclosure with simultaneous improvements in forecast agreement and accuracy (Lang and Lundholm 1996; Lehavy et al. 2011; Chircop et al. 2026); the analyst-information literature distinguishes common from private information and shows that analysts can herd or strategically compress observable disagreement (Barron et al. 1998; Clement and Tse 2005; Liu and Natarajan 2012); and the washing literature studies discrepancies between corporate claims and underlying outcomes (Kim and Lyon 2015; Baker et al. 2024). What these literatures do not establish is whether narrative–substance decoupling can break the conventional link between consensus and information quality. If analysts place similar weight on an AI narrative that is subsequently weakly substantiated, their forecast errors can become more correlated: analysts agree more because they share the same informational anchor, not because fundamentals have become more precisely understood. We refer to this configuration as false consensus—greater cross-analyst agreement without a corresponding improvement in the accuracy of the underlying belief. Accordingly:

H2: *Greater ex post AI narrative–output decoupling is associated with greater analyst forecast error.*

2.5 Analyst Information Production and Susceptibility to Common Narratives

Analysts are not passive recipients of corporate disclosure. They discover information independently, interpret public information, communicate with management, and draw on accumulated firm and industry knowledge (Barron et al. 2002b; Chen et al. 2010; Brown et al. 2015). These activities can reduce reliance on a common corporate narrative by providing alternative signals against which managerial claims can be evaluated. Barron, Byard, and Kim (2002b) show that public earnings information can stimulate additional idiosyncratic information production, and Barron, Byard, and Yu (2017) show that some disclosures increase private-information precision rather than simply common information. Recent evidence further indicates that information-processing capacity matters: Shanthikumar and Yoo (2026) find that AI resources allow sell-side analysts to process public information more efficiently and reallocate capacity toward private information acquisition, resulting in bolder and higher-quality forecasts.

The influence of difficult-to-verify common narratives should consequently vary with analysts' capacity to generate independent information. Firm-specific experience improves analysts' forecasting performance (Mikhail et al. 1997), while analyst ability, brokerage resources, and portfolio complexity affect forecast accuracy and information production (Clement 1999). Clement and Tse (2005) further show that prior accuracy, brokerage size, and experience are associated with a greater propensity to issue bold rather than herding forecasts, consistent with better-resourced analysts placing more weight on their own information. Brown et al. (2015) similarly emphasize the importance of industry knowledge and private communication in analysts' research process. Analysts with stronger firm-specific knowledge or greater information-production capacity should therefore be less dependent on a shared but difficult-to-verify AI narrative, limiting the extent to which narrative–output decoupling translates into common beliefs at the expense of differentiated private information.

3 Measurement and Empirical Design

3.1 Sample and Data Sources

Our sample consists of U.S. publicly listed firms over 2014–2023. We combine information from SEC EDGAR, Compustat, CRSP, I/B/E/S, and the United States Patent and Trademark Office (USPTO). The unit of observation in the baseline analyses is the firm–fiscal-year. Institutional ownership is obtained from Thomson Refinitiv 13F data, management earnings guidance from the First Call Company Issued Guidance database, and the Bog Index from Brian Miller's publicly available 10-K readability database.

The timing of the variables is important for our empirical design. We index each observation by the fiscal year whose earnings analysts are forecasting, denoted by (t). For each observation, AI-related narrative information is obtained from the Form 10-K corresponding to fiscal year *t-1*, which is normally filed during fiscal year *t*. We measure analyst disagreement using the first valid FY1 analyst consensus following that 10-K filing. The consensus therefore pertains to earnings for fiscal year *t*. Subsequently verifiable AI technological output is measured beginning immediately after the first post-10-K FY1 consensus date and through the end of fiscal year *t*. This sequencing ensures that the technological output used to classify the focal disclosure as AI washing is not publicly available when analysts form the consensus used in our primary dependent variable.[1] The post-filing consensus is constructed from analysts' latest outstanding forecasts as of the first valid post-filing consensus date. Some forecasts may have been issued before the filing and carried forward without revision. Accordingly, this measure represents the prevailing forecast distribution after the disclosure became public, rather than a sample composed exclusively of forecasts newly issued in response to the filing. The event-time and actual-revision analyses provide complementary evidence on changes in forecast distributions and analyst behavior around the disclosure.

This timing convention serves two purposes. First, the firm's AI-related narrative

[1] Because we index observations by the fiscal year being forecast, the first sample observation in 2014 generally relies on the fiscal-year 2013 Form 10-K filed during 2014.

is publicly available before the post-filing consensus snapshot used to measure analyst disagreement. Second, the technological output used to assess whether that narrative is subsequently substantiated is observed only after the analyst consensus used in our primary tests. Accordingly, *AIWashing* is an ex post classification of the extent to which the focal AI-related narrative is subsequently substantiated, rather than information available to analysts when they form their forecasts.

We obtain annual reports and filing dates from SEC EDGAR and extract the Management's Discussion and Analysis of Financial Condition and Results of Operations (MD&A) section from each original Form 10-K. We focus on MD&A because it contains management's discussion of operating conditions, investment priorities, technological initiatives, and expected developments and has been widely used to construct disclosure-based measures in accounting and finance research (Li 2008; Loughran and McDonald 2011; Bochkay et al. 2023).[2]

Firm-level accounting information is obtained from Compustat. These data provide assets, earnings, sales, R&D expenditures, segment information, and other accounting variables used to construct the control variables and supplementary measures. CRSP provides stock prices and returns.

Individual analysts' annual EPS forecasts and realized EPS are obtained from the I/B/E/S Detail History and Actuals files. Following Lehavy et al. (2011) and Fu et al. (2025), we construct the first post-10-K FY1 consensus from individual forecasts in the I/B/E/S Detail History file. We define the consensus date as the first valid post-filing consensus date for the subsequent fiscal period. For each analyst covering firm *i*, we retain the analyst's most recent outstanding FY1 EPS forecast as of that date, so that each analyst contributes no more than one forecast to the consensus.

AI-related technological output is identified using the USPTO Artificial Intelligence Patent Dataset 2023 (AIPD 2023; Pairolero et al. 2025). AIPD 2023 applies

[2] We use the original Form 10-K rather than a subsequently filed 10-K/A because an amendment may contain information that was unavailable to analysts when they initially responded to the annual filing.

a machine-learning classification procedure to U.S. patents and pre-grant publications to identify patent documents containing AI technologies. The dataset covers approximately 15.4 million patent documents published through 2023. Patent assignees are linked to public firms using historical assignee names, subsidiaries, ownership relationships, and corporate identifiers. The matching procedure follows the general principles of Arora et al. (2021), which emphasize the importance of historical name changes and ownership structures when linking patents to public firms.[3]

We link Compustat and CRSP through the CRSP/Compustat Merged database and match I/B/E/S forecasts to CRSP securities using historical identifier links. Patent assignees are subsequently matched to Compustat firms using historical corporate names and ownership relationships. All matches are validated to avoid duplicate firm-year observations.[4]

Our initial sample consists of U.S. firms with CRSP share codes 10 or 11 and exchange codes 1, 2, or 3, corresponding to ordinary common shares traded on the NYSE, AMEX, or NASDAQ. Following prior accounting research, we exclude financial firms (SIC 6000–6999) and utilities (SIC 4900–4999), as these industries are subject to distinct regulatory and financial-reporting requirements and exhibit systematically different disclosure characteristics (Fu et al. 2025). We further require nonmissing information necessary to construct AI washing, analyst forecast dispersion, and the baseline controls. Unless otherwise stated, continuous control variables are winsorized at the 1st and 99th percentiles. AI Narrative Intensity and Actual AI Output are not winsorized before percentile ranking because the rank transformation itself limits the influence of extreme observations. These screening procedures yield a final sample of 13,450 firm-year observations over the 2014–2023 period. Appendix Table A1 reports the sample-selection procedure, while Appendix Table A2 provides detailed

[3] Matching patent assignees only to contemporaneous firm names can create substantial measurement error following acquisitions, subsidiary reorganizations, and corporate name changes. Historical names and ownership relationships are therefore retained in the matching process.

[4] We distinguish firm-years with no observed patent publications from observations for which reliable patent-assignee linkage cannot be established. The former are retained and assigned zero patent-based AI output, whereas the latter are excluded from the analysis.

variable definitions and data sources.

3.2 Variable Construction

3.2.1 Explanatory Variable

Following the broader corporate-washing literature, we conceptualize AI washing as a discrepancy between the intensity of a firm's external AI-oriented communication and its subsequently verifiable substantive AI technological activity (Kim and Lyon 2015; Baker et al. 2024). This definition focuses on narrative–output decoupling rather than on a low level of AI activity per se. Accordingly, a firm is characterized as exhibiting greater AI washing when its relative emphasis on AI in corporate disclosure exceeds its relative subsequently verifiable AI technological output.

We first measure firms' AI-oriented narratives using the MD&A section of the focal Form 10-K. Following established textual-analysis practice (Loughran and McDonald 2011; Bochkay et al. 2023), we construct a domain-specific AI dictionary covering core AI technologies, including artificial intelligence, machine learning, deep learning, neural networks, natural language processing, computer vision, reinforcement learning, and related technologies.[5] The dictionary draws on the OECD AI classification framework, the ACM Computing Classification System, the technological categories underlying the USPTO Artificial Intelligence Patent Dataset (AIPD), and prior research measuring corporate AI disclosures (Ante and Saggu 2025;

[5] AI-related expressions are identified using case-insensitive phrase matching, except for standalone abbreviations, for which capitalization rules are applied. Longer expressions are matched first to prevent overlapping counts. Multiword expressions contribute their constituent word counts to the numerator, consistent with the word-based denominator. We additionally screen ambiguous abbreviations to reduce false-positive matches.

Barrios et al. 2026).[6] Generic digitalization terms such as "*digital*, *automation*, *smart*, and *information technology*" are excluded unless they form part of an explicitly AI-related expression. We define AI Narrative Intensity as:

$$AI\ Narrative\ Intensity_{i,t} = \frac{AIWords_{i,t}}{MD\&AWords_{i,t}}, \quad (1)$$

where $AIWords_{i,t}$ is the number of AI-related words identified in the focal MD&A and $MD\&AWords_{i,t}$ is the total number of words in that section. A higher value therefore indicates that AI-related technologies occupy a greater share of the firm's narrative discussion. Because the focal 10-K is publicly available before the first post-filing analyst consensus used in our main tests, this component of *AIWashing* belongs to analysts' information set when they form their forecasts. This timing is consistent with the sequencing described in Section 3.1.

We next measure Actual AI Output using AIPD 2023 (Pairolero et al. 2025). Specifically, Actual AI Output is the proportion of a firm's pre-grant patent publications classified by AIPD as containing AI technologies among all of its pre-grant patent publications during the verification window:

$$ActualAIOutput_{i,t} = \frac{AI\ PGPubs_{i,t}}{All\ PGPubs_{i,t}}, \quad (2)$$

The verification window begins on the day following the first valid post-10-K FY1 consensus date and ends at the end of fiscal year *t*.[7] We use pre-grant publications because they provide an earlier publicly observable patent-based signal than granted patents (Pairolero et al. 2025). For firm-years with no pre-grant patent publications

[6] The dictionary is restricted to identifiable AI technologies rather than broad digitalization terminology. Appendix Table A3 reports the complete keyword list and coding procedure.

[7] AIPD 2023 classifies publicly available patent documents. We therefore interpret Actual AI Output as publicly verifiable patent-based technological output rather than the universe of confidential patent applications.

during the verification window, Actual AI Output is set to zero,[8] reflecting the absence of publicly observable patent-based AI output during that period. Appendix Table A7 further examines the positive-patenting subsample and an alternative output measure based on the log number of AI pre-grant patent publications. Importantly, because AIPD captures patent-based technological activity rather than the universe of firms' AI implementation, Actual AI Output should be interpreted as subsequently verifiable patent-based AI output, rather than as a comprehensive measure of all AI adoption or investment.[9]

Because AI Narrative Intensity and Actual AI Output are measured on different natural scales, directly subtracting their raw values would make the resulting gap depend mechanically on the relative scaling and distribution of the two components. We therefore adapt the within-year percentile-rank approach of Baker et al. (2024). For each forecast year, we independently rank AI Narrative Intensity and Actual AI Output and normalize the resulting percentile ranks to the $[0,1]$ interval. Average ranks are assigned to tied observations. We then define *AIWashing* as:

$$AIWashing_{i,t} = PctRank_t\left(AI\ Narrative\ Intensity_{i,t}\right) - PctRank_t\left(Actual\ AI\ Output_{i,t}\right), \quad (3)$$

[8] Coding nonpatenting firm-years as zero reflects the absence of observable patent-based AI output rather than the mathematical value of an otherwise undefined ratio. The positive-patenting subsample test assesses the sensitivity of the results to this treatment.

[9] Patent publications provide observable evidence of AI-related technological activity, but they do not capture the full scope of corporate AI implementation. Firms may adopt AI through externally developed software, licensed technologies, or third-party services without generating proprietary AI patents. Prior research accordingly measures firm-level AI investment using employee skills and recruitment activity (Babina et al. 2024), while survey evidence distinguishes actual AI adoption from technological invention (McElheran et al. 2024). Moreover, patent publication typically follows the underlying application, meaning that a patent published during the verification window need not represent innovation initiated after the focal disclosure. We therefore interpret Actual AI Output as subsequently observable patent-based technological evidence, rather than a comprehensive measure of AI adoption or newly undertaken innovation.

where $PctRank_t$ denotes the percentile rank of a firm within forecast year t.[10] The interpretation of *AIWashing* is therefore relative rather than absolute.[11] A positive *AIWashing* value indicates that a firm's relative AI narrative intensity exceeds its subsequently verifiable AI technological output, whereas a negative value indicates comparatively stronger realized AI output. Because both components are transformed into within-year percentile ranks, the mean of *AIWashing* is approximately zero by construction; thus, its economic meaning lies in the sign and magnitude of the narrative–output gap rather than in the unconditional mean or a sharp zero threshold.[12]

3.2.2 Dependent Variable

Our primary dependent variable is analyst forecast disagreement. Following Lehavy et al. (2011) and Fu et al. (2025), we construct the first post-10-K FY1 consensus for the subsequent fiscal period using individual forecasts from I/B/E/S Detail History. For each analyst covering firm *i*, we retain the most recent outstanding FY1 EPS forecast as of the consensus date and require forecasts from at least two distinct analysts to calculate dispersion:

$$FDISP_{i,t} = \frac{SD_j(FEPS_{j,i,t})}{Price_{i,t}^{-90}}, \quad (4)$$

where $FEPS_{i,j,t}$ denotes analyst j's most recent outstanding FY1 EPS forecast for firm i at the first valid post-10-K consensus date in forecast year t, and $SD_j(\cdot)$

[10] Baker et al. (2024) use the same basic logic in their diversity-washing measure: the disclosure component and underlying substantive outcome are ranked separately within year and the two ranks are differenced. They motivate this construction as a way to compare differently scaled measures, mitigate sensitivity to extreme observations and measurement error, and absorb common temporal shifts in disclosure.

[11] *AIWashing* is a relative measure of ex post narrative–patent-output decoupling. A higher value indicates that a firm's AI narrative occupies a stronger position in the annual cross-sectional distribution than its subsequently observable patent-based AI output. The measure does not establish intentional exaggeration, nor does it distinguish deceptive reporting from genuine AI adoption that generates no proprietary patents. Accordingly, we use AI washing as an operational label for narrative–output decoupling, rather than as a direct measure of managerial misrepresentation.

[12] Importantly, AI Narrative Intensity is observable from the focal 10-K when analysts form their forecasts, whereas Actual AI Output is measured only after the first post-10-K FY1 consensus. *AIWashing* therefore represents an ex post assessment of whether the focal AI narrative is subsequently substantiated by publicly verifiable technological output, allowing us to examine whether disclosures later characterized by larger narrative–output gaps are associated with greater contemporaneous analyst convergence.

denotes the cross-sectional standard deviation across analysts covering the same firm. $Price_{i,t}^{-90}$ is the absolute split-adjusted CRSP share price 90 calendar days before the consensus date, using the most recent preceding trading-day price when necessary. We require at least two distinct analysts to construct $FDISP$. Following Lehavy et al. (2011) and Fu et al. (2025), we scale forecast dispersion by stock price rather than consensus EPS to ensure comparability across firms and avoid denominator instability when expected earnings approach zero. Higher FDISP indicates greater disagreement among analysts.

3.2.3 Control Variables

Following prior research on analyst forecast properties, we control for firm characteristics, business complexity, and the information environment (Lehavy et al. 2011; Park and Zach 2025). *SIZE* is the natural logarithm of market capitalization, in millions of U.S. dollars, at fiscal-year end *t-1*. *GROWTH* is the compound annual growth rate of sales over the preceding five fiscal years; when five years of history are unavailable, we use the longest available window of at least three years, requiring positive beginning and ending sales. *SEGMENTS* is the natural logarithm of the number of reported business segments. *INST* is the fraction of shares outstanding held by institutional investors based on the latest publicly available 13F holdings information preceding the focal 10-K filing. *MFCOUNT* is the number of firm-issued earnings-guidance announcements during the preceding 12 months, with duplicate guidance for the same fiscal period on the same date counted once. *STDRET* is the standard deviation of monthly stock returns during fiscal year *t-1*. *10KNEWS* is the absolute cumulative market-adjusted return over trading days [0,+1] around the focal 10-K filing, using the CRSP value-weighted market return as the benchmark. For filings submitted after market close or on a non-trading day, day 0 is the next trading day.

We additionally control for fundamental uncertainty and the disclosure environment. *EARNVOL* is the standard deviation of quarterly earnings before extraordinary items scaled by quarterly total assets (IBQ/ATQ) over the eight quarters

preceding the focal filing. *LOSS* equals one if annual income before extraordinary items (IB) is negative in fiscal year *t-1*, and zero otherwise. *FOLLOW* is the natural logarithm of one plus the number of distinct analysts with an outstanding FY1 EPS forecast immediately before the focal 10-K filing. *FHORIZON* is the natural logarithm of one plus the number of calendar days between the initial post-filing FY1 consensus date and target fiscal-year end. *BOG* is the Bog Index of the focal original 10-K (Bonsall et al. 2017; Fu et al. 2025). Finally, *RDINT* equals R&D expenditure divided by sales in fiscal year *t-1*, with missing R&D expenditure set to zero and observations with missing or nonpositive sales excluded. For a detailed description of the variables, please refer to Table A2 in the Appendix.

3.3 Empirical Specification

To examine whether analysts exhibit greater consensus following corporate disclosures that are subsequently characterized by greater AI washing, we estimate the following baseline model:

$$FDISP_{i,t} = \alpha + \beta_1 AIWashing_{i,t} + \gamma' Controls_{i,t} + FirmFE_i + YearFE_t + \varepsilon_{i,t}, \quad (5)$$

where $FDISP_{i,t}$ analyst forecast disagreement in the first FY1 consensus following firm i's focal 10-K filing in year t, and $AIWashing_{i,t}$ is ex post narrative–output gap; $Controls_{i,t}$ is the vector of control variables. This construction follows prior studies that measure analyst forecast properties using the first consensus following the annual filing for the subsequent fiscal period (Lehavy et al. 2011; Fu et al. 2025). The Firm fixed effects ($FirmFE$) absorb time-invariant firm characteristics, while fiscal-year fixed effects ($YearFE$) control for common changes in analyst behavior, disclosure practices, and the salience of AI. The coefficient of interest is β_1. A negative β_1 indicates that larger narrative–output rank gaps are associated with lower analyst forecast disagreement following the focal 10-K filing. Because *AIWashing* is an ex post percentile-rank difference rather than a contemporaneously observable treatment, we assess its economic magnitude using a one-standard-deviation change rather than its

unconditional mean and interpret the estimates as conditional associations rather than definitive causal effects.

4 Main Results

4.1 Descriptive Statistics and AI-Washing Validation

Table 1 reports descriptive statistics for 13,450 firm-year observations. *AIWashing* has a mean of zero by construction and a standard deviation of 0.4042, indicating considerable variation in narrative–output discrepancies. Panel B shows that firms with higher *AIWashing* exhibit significantly lower analyst forecast disagreement (0.0132 versus 0.0139) but higher forecast errors (0.0269 versus 0.0253).[13] Both differences are significant at the 1% level. This contrasting pattern provides preliminary evidence consistent with false consensus: greater agreement among analysts does not necessarily translate into more accurate forecasts.

<Insert Table 1 here>

Figure 1 illustrates the joint distribution of AI Narrative Intensity and subsequently verifiable AI Output. Two features are noteworthy. First, although higher narrative intensity is generally associated with stronger patent-based AI output, the relationship is far from one-to-one. In particular, a substantial number of firm-years in the higher narrative groups remain concentrated in the Zero AI Output category. Second, even among observations with positive AI output, firms with similar levels of narrative intensity are distributed across multiple output groups. Taken together, these patterns indicate considerable heterogeneity in the alignment between AI-oriented disclosure and subsequently verifiable technological activity. Appendix Tables A4 and A5 provide complementary descriptive evidence, while Appendix Figure A1 documents the distribution and annual incidence of *AIWashing*.

<Insert Figure 1 here>

[13] Observations equal to the sample median are assigned to the Low AIWashing group, resulting in slightly unequal group sizes.

4.2 AI Washing and Analyst Forecast Disagreement

Table 2 reports the baseline regression results. Column (1) includes year fixed effects, Column (2) further incorporates firm fixed effects, and Column (3) replaces year fixed effects with industry-by-year fixed effects. Focusing on Column (2), which includes firm and year fixed effects, the coefficient ($\beta = -0.0010$) on *AIWashing* is significantly negative at the 1% level. Economically, a one-standard-deviation increase in *AIWashing* is associated with lower analyst forecast disagreement, with an estimated difference equivalent to 2.97% of the sample mean.[14] This finding suggests that greater narrative–output decoupling is associated with stronger apparent consensus among financial analysts.

<Insert Table 2 here>

4.3 Event-Time Evidence

To examine the timing of analyst convergence, we construct an event-time panel around each focal 10-K filing using six trading-day windows: [−60,−31], [−30,−11], [−10,−1], [0,+10], [+11,+30], and [+31,+60]. At the end of each window, we retain each analyst's most recent outstanding FY1 forecast and construct price-scaled cross-sectional forecast dispersion using the same price denominator as in the baseline analysis. A firm-year-window is retained only when at least two distinct analysts have valid outstanding forecasts. Post equals one for the three post-filing windows and zero for the three pre-filing windows. We estimate interactions between *AIWashing* and the event-time indicators with firm-year and event-window fixed effects.

Figure 2 shows that the association between AIWashing and analyst forecast dispersion is weak before the focal 10-K filing but becomes negative afterward. This temporal pattern is consistent with stronger post-filing convergence among firms subsequently characterized by larger narrative–output gaps. However, because the

[14] Economic significance is calculated as $|\beta| \times SD(AIWashing)/Mean(FDISP) \times 100$. Using the coefficient from Column (2) and the sample statistics reported in Table 1, the estimated reduction is 0.0010 × 0.4042 / 0.0136 × 100 = 2.97% (based on the reported rounded coefficient).

filing contains financial and nonfinancial information beyond AI-related disclosures, the event-time estimates do not isolate the causal effect of AI narratives. The evidence is therefore interpreted as supporting the timing of the documented association rather than establishing a disclosure-specific causal mechanism.

<Insert Figure 2 here>

Table 3 formally examines this relationship. Column (1) reports the full-window specification, Column (2) restricts the event window to ±30 trading days, and Column (3) retains a stable set of analysts. Focusing on Column (1), the coefficient on *AIWashing* × *Post* is significantly negative at the 1% level. Economically, a one-standard-deviation increase in *AIWashing* is associated with an additional post-filing difference in forecast dispersion equivalent to 0.59% of the sample mean.[15] The negative association remains consistent across alternative specifications, providing complementary evidence of stronger post-filing analyst convergence among firms subsequently characterized by greater narrative–output gaps.

<Insert Table 3 here>

4.4 Complementary Instrumental-Variable Evidence

To complement the baseline analysis, we adopt an instrumental-variable approach following the historical firm–university network design of Babina et al. (2024). We construct each firm's pre-sample exposure to AI-strong universities using the educational backgrounds of its STEM workforce in 2010 and universities' AI research activity during 2005–2009. A university is classified as AI-strong if, in at least one year during 2005–2009, it ranks either in the top 5% by the number of AI researchers or simultaneously in the top 10% by the number and top 5% by the share of AI researchers.

[15] Economic significance is calculated as $|\beta| \times SD(AIWashing)/Mean(FDISP) \times 100 = 0.0002 \times 0.4042/0.0136 \times 100 = 0.59$, using the rounded coefficient in Column (1) and the full-sample descriptive statistics in Table 1 as the benchmark.

AIExposure is the STEM-worker-weighted share of a firm's university connections attributable to AI-strong universities, and *LowAIExposure (= 1 − AIExposure)* serves as the excluded instrument. Using the same firm–university weights, we control for exposure to 2010 U.S. News top-10 universities and universities' historical non-AI computer-science research strength. All measures are fixed at their pre-sample values. Because historical talent networks may affect innovation and firms' information environments through channels other than *AIWashing*, the resulting IV estimates are interpreted as complementary evidence conditional on the exclusion restriction rather than definitive causal identification.[16]

We implement a two-stage least squares estimation. In the first stage, we estimate the following model:

$$AIwashing_{i,t} = \alpha_0 + \alpha_1 LowAIExposure_i + \delta_1 TopUnivExposure_i + \delta_2 CSUnivExposure_i + \gamma Controls_{i,t} + \lambda_j + \theta_t + \varepsilon_{i,t}\,, \quad (6)$$

where *LowAIExposure* is defined as one minus a firm's historical exposure to AI-strong universities, calculated using the educational backgrounds of its STEM workforce in 2010. *TopUniExposure* and *CSStrongExposure* are constructed using the same firm–university weights, replacing AI research strength with overall university standing and non-AI computer-science research strength, respectively. All measures are fixed at their

[16] The validity of this instrument requires that historical AI-talent exposure affects analyst forecast disagreement only through AIWashing, conditional on the included covariates and fixed effects. This restriction is difficult to establish in our setting. Historical AI-talent networks may directly influence firms' innovation capabilities, patent output, operating performance, and disclosure practices, all of which may affect analysts' information environments. Moreover, because patent-based output is a component of AIWashing, instrument relevance may partly reflect differences in technological production rather than variation specific to narrative–output decoupling. Although we control for university quality, non-AI computer science research strength, and R&D intensity, these controls do not eliminate the possibility of direct effects. The IV analysis is therefore treated as supplementary evidence, without attributing a causal interpretation to the second-stage estimates (Larcker and Rusticus 2010; Gow et al. 2016).

pre-sample values. Industry and year fixed effects are included, rather than firm fixed effects, because the excluded instrument is time-invariant.[17] A positive first-stage coefficient would indicate that weaker historical AI-talent exposure is associated with larger narrative–output gaps, conditional on the included covariates and fixed effects.

In the second stage, we estimate the relationship between instrumented *AIWashing* and *FDISP* using two-stage least squares:

$$FDISP_{i,t+1} = \beta_0 + \beta_1 \widehat{AIwashing}_{i,t} + \delta_1 TopUnivExposure_i + \delta_2 CSUnivExposure_i + \gamma Controls_{i,t} + \lambda_j + \theta_t + \varepsilon_{i,t+1}, \quad (7)$$

where $\widehat{AIwashing}$ denotes the fitted value of $AIwashing$ obtained from the first-stage regression using $LowAIExposure$ as the instrumental variable. A negative second-stage coefficient would be consistent with the baseline association between greater AI washing and lower analyst forecast disagreement, conditional on instrument validity. Given the maintained exclusion restriction and the replacement of firm fixed effects with industry fixed effects, we interpret the IV results as supplementary evidence rather than conclusive causal identification.

Table 4 reports supplementary 2SLS estimates. LowAIExposure is positively associated with AIWashing in the first stage[18], with a Kleibergen–Paap rk Wald F-statistic of 42.92. The second-stage coefficient is negative, consistent in direction with the baseline association. However, first-stage strength and weak-instrument-robust

[17] Since *LowAIExposure* is a pre-sample, time-invariant firm-level measure, it cannot be separately identified in a specification with firm fixed effects. The IV regressions therefore replace firm fixed effects with industry fixed effects and year fixed effects, while retaining the baseline controls.

[18] The IV sample comprises 10,286 firm-year observations for which firms' pre-sample STEM employee–university networks and universities' historical research characteristics can be reliably identified. The reduction from the baseline sample of 13,450 observations reflects the additional data requirements for constructing the historical AI-talent exposure instrument. Firms without reliable historical network information are excluded rather than assigned zero exposure.

inference do not establish the exclusion restriction. In addition, the IV specification relies on a smaller sample and industry rather than firm fixed effects, limiting comparability with the baseline estimates. We therefore interpret these results cautiously as supplementary evidence. Appendix Table A13 reports additional instrument diagnostics and weak-instrument-robust inference. The underidentification test rejects the null, and the Anderson–Rubin test remains significant at the 5% level, providing complementary evidence on instrument relevance and inference under potential weak identification.

<Insert Table 4 here>

4.5 AI Washing and Forecast Accuracy

We measure *Consensus Forecast Error* as the absolute difference between the first post-10-K FY1 consensus EPS forecast and subsequently realized EPS, scaled by the same price denominator used for *FDISP*. *Analyst Forecast Error* is constructed analogously using each analyst's outstanding FY1 EPS forecast. Realized earnings are obtained from I/B/E/S Actuals, with forecasts and actual EPS maintained on a consistent split-adjusted basis.

Table 5 examines whether greater analyst consensus is accompanied by improved forecast accuracy. Column (1) examines consensus forecast error, while Column (2) considers individual analyst forecast error. Focusing on Column (1), the coefficient on *AIWashing* is positive and significant at the 1% level. The positive association also holds at the individual analyst level, suggesting that the lower forecast disagreement documented earlier does not translate into greater forecast accuracy. This contrasting pattern is consistent with our interpretation of false consensus.

<Insert Table 5 here>

4.6 Evidence Consistent with False Consensus

Figure 3 illustrates the contrasting associations between *AIWashing* and analyst forecast properties. Panel A shows a declining pattern in forecast disagreement across *AIWashing* deciles, whereas Panel B exhibits an increasing pattern in forecast error. Together, these patterns provide visual evidence consistent with false consensus, whereby greater analyst agreement is accompanied by less accurate earnings expectations.

<Insert Figure 3 here>

To examine subsequent forecast adjustments, we compare the initial post-10-K FY1 consensus with the last valid consensus preceding the corresponding earnings announcement. *Revision Magnitude* is defined as the absolute difference between these two consensus forecasts, scaled by the same stock-price denominator used for *FDISP*. *Revision toward Realized EPS* equals one if the absolute forecast error of the later consensus is smaller than that of the initial consensus, and zero otherwise. We obtain forecast histories and realized EPS from I/B/E/S, with earnings-announcement dates identified using I/B/E/S Actuals.[19]

Table 6 further examines subsequent forecast adjustments. Column (1) considers the magnitude of consensus forecast revisions, while Column (2) examines whether revisions move toward realized EPS. Focusing on Column (1), *AIWashing* is positively associated with subsequent revision magnitude, significant at the 1% level. Column (2) further indicates a positive association with corrections toward realized earnings. These findings suggest that the tighter initial consensus associated with greater AI washing is also followed by more pronounced forecast adjustments, consistent with the interpretation that initial agreement does not necessarily reflect superior information.

<Insert Table 6 here>

5 Analyst Behavior and Information-Environment Evidence

[19] We require the later consensus to be distinct from the initial post-filing consensus and to refer to the same target fiscal period; firm-years without both valid consensus snapshots are excluded from this analysis.

5.1 Analyst-Level Convergence toward the Prevailing Consensus

To examine individual analysts' convergence toward the prevailing consensus, we construct analyst-level distance measures following the leave-one-out approach. Specifically, *Distance to Consensus* is calculated as the absolute difference between an analyst's outstanding FY1 EPS forecast and the mean forecast of all other analysts covering the same firm, scaled by the same price denominator used for *FDISP*. *Pre-Filing Distance* and *Post-Filing Distance* are measured using the last valid pre-filing and first valid post-filing consensus, respectively, retaining analysts with valid forecasts in both periods. *Change in Distance* is defined as Post-Filing Distance minus Pre-Filing Distance, with negative values indicating convergence.[20]

Table 7 reports the results. Column (1) examines *Post-Filing Distance* while controlling for *Pre-Filing Distance*, whereas Column (2) uses *Change in Distance* as the dependent variable. Focusing on Column (1), the coefficient on *AIWashing* is negative and significant at the 1% level. The negative association persists in Column (2), suggesting stronger post-filing convergence among analysts covering firms subsequently characterized by greater AI washing.

<Insert Table 7 here>

5.2 Analyst Herding and Forecast Boldness

Following Clement and Tse (2005), we classify each analyst's first actual FY1 forecast revision after the focal 10-K filing and before the corresponding earnings announcement relative to the analyst's immediately preceding forecast and the prevailing pre-revision leave-one-out consensus. *HERD* equals one when the revised forecast lies strictly between these benchmarks, whereas *BOLD* equals one when it lies strictly outside their interval. Revisions equal to either benchmark are retained and coded as zero for both indicators. Analyst-firm-years without an identifiable post-filing

[20] Both distance measures use the same price denominator and are calculated for analysts identifiable in both the pre- and post-filing forecast snapshots. Analyst identities and forecast histories are obtained from I/B/E/S Detail History. The leave-one-out consensus is calculated separately for each snapshot, excluding the focal analyst's own forecast.

revision are excluded.

Table 8 reports the results using linear probability models. Column (1) shows that *AIWashing* is positively associated with *HERD*, significant at the 5% level, whereas Column (2) documents a negative association with *BOLD*, significant at the 10% level. These findings are consistent with greater herding and less bold forecasting among analysts covering firms subsequently characterized by higher AI washing, providing complementary behavioral evidence of apparent consensus.

<Insert Table 8 here>

5.3 Model-Implied Analyst Information Composition

To further examine analysts' information environments, we follow Barron et al. (1998) and decompose forecast information into common and private components. Let D denote the cross-sectional variance of price-scaled individual forecasts, SE the squared price-scaled consensus forecast error, and N the number of analysts. Common uncertainty is estimated as $C = SE - D/N$, and total uncertainty as $V = C + D$. We subsequently construct consensus ($\rho = C/V$), common information precision ($h = C/V^2$), and private information precision ($s = D/V^2$).[21] We retain only observations for which the model-implied common uncertainty C and total uncertainty V are strictly positive and all three information measures are finite and defined. Consensus, common information precision, and private information precision are subsequently standardized to mean zero and unit standard deviation within this valid sample.

Table 9 reports model-implied measures of analysts' information composition following Barron et al. (1998). AIWashing is positively associated with the consensus measure and negatively associated with estimated private-information precision, while the coefficient on common-information precision is statistically insignificant. Because these measures are constructed from forecast dispersion and consensus forecast error,

[21] All components are calculated using the same post-filing consensus, analyst coverage count, and price-scaling convention. Because the decomposition relies on model-implied information parameters, we separately examine the validity of the constructed measures, including observations with nonpositive estimated uncertainty or undefined precision measures. The resulting variables are interpreted as model-implied proxies for analysts' information composition rather than direct observations of private information acquisition.

they are not independent of the forecast outcomes examined earlier. We therefore interpret the estimates as model-based descriptions of the information environment, rather than direct evidence of changes in analysts' private-information acquisition.

<Insert Table 9 here>

5.4 Heterogeneity by Analyst Experience

To examine whether analyst experience moderates the association between AI washing and forecast convergence, we define *High Firm Experience* as an indicator equal to one if an analyst's firm-specific coverage tenure, measured as the elapsed time between the analyst's first recorded FY1 forecast for the focal firm and the 10-K filing date, exceeds the sample-year median, and zero otherwise. Coverage histories are obtained from I/B/E/S Detail History, including available pre-sample records to reduce left-censoring of analyst experience.

Table 10 reports the interaction results. Column (1) examines *Change in Distance*, while Column (2) considers *HERD*. The interaction between *AIWashing* and *High Firm Experience* is positive in Column (1) and negative in Column (2), both significant at the 5% level. These findings suggest that the associations between greater AI washing, forecast convergence, and herding are less pronounced among analysts with greater firm-specific experience. Appendix Table A12 reports complementary heterogeneity analyses based on analysts' prior forecast accuracy, brokerage size, and industry specialization.

<Insert Table 10 here>

6 Supplemental Tests

6.1 AI Narrative–Patent-Output Alignment

To distinguish AI narrative–output decoupling from observable patent-based AI activity, we classify firm-years into four mutually exclusive groups using a within-year median split of AI Narrative Intensity and an indicator for positive AI-related pre-grant patent output. *High Narrative* equals one when AI Narrative Intensity is strictly above

the within-year median, with observations at the median assigned to the low-narrative group. *Positive AI Output* equals one when at least one AI-classified pre-grant patent publication is observed during the verification window. This classification accommodates the substantial concentration of zero observations in Actual AI Output. *Low Narrative–Zero AI Output* serves as the reference group, while *High Narrative–Zero AI Output*, *High Narrative–Positive AI Output*, and *Low Narrative–Positive AI Output* constitute the remaining groups. Table 11 reports the results. Firms with high narrative intensity but zero observed AI output exhibit lower forecast disagreement, higher forecast error, and larger subsequent forecast revisions, whereas those with high narrative intensity and positive AI output exhibit both lower disagreement and lower forecast error. These contrasting patterns suggest that the informational content of analyst consensus differs with observable technological support for corporate AI narratives. However, because patent-based output does not capture all forms of AI adoption, these classifications do not establish whether firms engage in deceptive reporting or genuine AI implementation.

<Insert Table 11 here>

6.2 AI Washing and Subsequent Earnings-Announcement Reactions

To examine the capital-market implications of apparent consensus, we measure $|CAR[-1,+1]|$as the absolute cumulative market-adjusted return over trading days $[-1,+1]$ surrounding the subsequent earnings-announcement date, where market-adjusted return equals the firm's daily CRSP return minus the CRSP value-weighted market return.

Table 12 reports the results. Column (1) examines the association between *AIWashing* and absolute announcement returns, while Column (2) additionally controls for *ForecastError*. Focusing on Column (1), the coefficient on *AIWashing* is positive and significant at the 5% level. The association remains significant after controlling for forecast error in Column (2). These results provide complementary market-based evidence that lower analyst disagreement is not necessarily associated with smaller

subsequent earnings-announcement reactions. The estimates do not establish that narrative–output decoupling causes greater market uncertainty or that forecast consensus mediates the announcement response. Appendix Table A14 reports complementary market-based outcomes using a wider announcement window, abnormal trading volume, and post-announcement absolute drift. The results are directionally consistent with the main announcement-return evidence.

<Insert Table 12 here>

6.3 Additional tests

We conduct several additional analyses to assess the robustness and scope of the main findings. First, we examine alternative constructions of both AI washing and analyst disagreement. Appendix Table A6 shows that the results are robust to a commitment-based narrative–output gap, industry-year ranking, and a binary AI-washer classification. Appendix Table A7 further considers alternative AI-output measures and verification windows, including the positive-patenting subsample, the log number of AI pre-grant publications, and fixed 180- and 365-day verification windows. Appendix Table A8 replaces the baseline forecast-dispersion measure with the interquartile range and median absolute deviation and imposes stricter minimum analyst-coverage requirements. Across these alternatives, greater AI washing remains associated with lower analyst forecast disagreement, while alternative AI-washing measures continue to exhibit a positive association with forecast error.

Second, we examine whether the baseline association is sensitive to alternative specifications, sample restrictions, or timing assumptions. Appendix Table A9 shows that the negative association between *AIWashing* and *FDISP* persists after introducing state-by-year fixed effects, excluding 2020, excluding the largest one percent of firms, requiring at least three analysts, and restricting the sample to patenting firm-years. Appendix Table A10 reports falsification tests using pre-filing disagreement, pre-filing changes, and a placebo lead outcome. None of these pre-filing or placebo coefficients is statistically significant, whereas the post-filing specification reproduces the negative

baseline association. Consistent with these results, Appendix Figure A2 extends the event-time analysis over a broader window and shows no comparable decline in disagreement before the focal 10-K filing.

Finally, we examine whether the documented patterns are concentrated in particular time periods or analyst types. Appendix Table A11 interacts *AIWashing* with an indicator for forecast-year 2023, the first forecast year following the public release of ChatGPT. The interaction is statistically insignificant for both forecast disagreement and forecast error, suggesting that the main associations are not concentrated exclusively in the first post-ChatGPT year. Appendix Table A12 considers additional analyst characteristics. The association between *AIWashing* and convergence toward the prevailing consensus is weaker among analysts with greater prior forecast accuracy, analysts employed by larger brokerages, and industry specialists. Together, these supplementary analyses indicate that the main findings are not driven by a particular measure, sample restriction, period, or analyst subgroup.

7 Conclusions

This study examines whether corporate AI narratives that are subsequently weakly substantiated by observable technological output shape analysts' information environment. Using 13,450 U.S. firm-year observations from 2014–2023, we measure AI washing as the within-year rank gap between AI narrative intensity in the focal 10-K and subsequently verifiable patent-based AI output. Because the output used to classify the disclosure is observed only after the initial post-filing consensus, AIWashing is an ex post characterization rather than information available to analysts when forecasts are formed. A one-standard-deviation increase in AIWashing is associated with a 2.97% reduction in forecast disagreement relative to its sample mean, but a 3.10% increase in consensus forecast error. The convergence emerges after the 10-K filing rather than before it and is followed by larger forecast revisions toward realized earnings. These findings are consistent with false consensus: analysts agree more closely without becoming more accurate.

The analyst-level evidence clarifies the information process behind this pattern. Greater AIWashing is associated with stronger convergence toward the prevailing consensus, more herding, less bold forecasting, and lower model-implied private-information precision, consistent with analysts placing greater relative weight on a common but difficult-to-verify corporate narrative and generating less differentiated forecasts. This association is weaker among analysts with greater firm-specific experience and, in supplementary analyses, among analysts with higher prior accuracy, larger brokerage resources, and greater industry specialization, suggesting that stronger information-production capacity mitigates reliance on the common narrative. Importantly, the pattern is distinct from substantive AI activity: strong AI narratives accompanied by observable AI output are associated with both lower disagreement and lower forecast error, whereas narrative–output decoupling is associated with lower disagreement but higher error.

Our findings suggest that analyst agreement and forecast accuracy capture distinct dimensions of the information environment. Greater consensus may coexist with less accurate expectations when firms' technology-oriented narratives are difficult to substantiate with observable evidence. Two qualifications are important. Patent publications capture only one dimension of AI activity and cannot establish whether managerial claims are misleading. Moreover, the empirical evidence identifies conditional associations rather than causal effects of AI washing on analyst behavior. Subject to these limitations, the results highlight why forecast dispersion alone may be an incomplete indicator of information quality when evaluating corporate disclosures about emerging technologies.

References

Ante, L., & Saggu, A. (2025). Quantifying a firm's AI engagement: Constructing objective, data-driven, AI stock indices using 10-K filings. *Technological Forecasting and Social Change*, *212*, 123965.

Arora, A., Belenzon, S., & Sheer, L. (2021). Matching patents to compustat firms, 1980–2015: Dynamic reassignment, name changes, and ownership structures. *Research Policy*, *50*(5), 104217.

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

Baker, A. C., Larcker, D. F., McClure, C. G., Saraph, D., & Watts, E. M. (2024). Diversity washing. *Journal of Accounting Research*, *62*(5), 1661-1709.

Barrios JM, Campbell JL, Johnson RG, Liu YC (2026) Signals or smoke? The determinants and informativeness of corporate artificial intelligence (AI) disclosures. *Review of Accounting Studies*, forthcoming.

Barron, O. E., Byard, D., & Kim, O. (2002b). Changes in analysts' information around earnings announcements. *The Accounting Review, 77*(4), 821–846.

Barron, O. E., Byard, D., & Yu, Y. (2017). Earnings announcement disclosures and changes in analysts' information. *Contemporary Accounting Research, 34*(1), 343–373.

Barron, O. E., Byard, D., Kile, C., & Riedl, E. J. (2002a). High-technology intangibles and analysts' forecasts. *Journal of Accounting Research, 40*(2), 289–312.

Barron, O. E., Kim, O., Lim, S. C., & Stevens, D. E. (1998). Using analysts' forecasts to measure properties of analysts' information environment. *The Accounting Review, 73*(4), 421–433.

Beyer, A., Cohen, D. A., Lys, T. Z., & Walther, B. R. (2010). The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics, 50*(2–3), 296–343.

Bhimani, A., Cinquini, L., & Malmi, T. (2026). What happens at the interface of digitalisation and accounting? The British Accounting Review, 58(2), 101742.

Bochkay, K., Brown, S. V., Leone, A. J., & Tucker, J. W. (2023). Textual analysis in accounting: What's next?. *Contemporary Accounting Research*, *40*(2), 765-805.

Bonsall, S. B., Leone, A. J., Miller, B. P., & Rennekamp, K. (2017). A plain English measure of financial reporting readability. *Journal of Accounting and Economics*, 63(2–3), 329–357.

Brown, L. D., Call, A. C., Clement, M. B., & Sharp, N. Y. (2015). Inside the "black box" of sell-side financial analysts. *Journal of Accounting Research, 53*(1), 1–47.

Chen, X., Cheng, Q., & Lo, K. (2010). On the relationship between analyst reports and

corporate disclosures: Exploring the roles of information discovery and interpretation. *Journal of Accounting and Economics, 49*(3), 206–226.

Chircop, J., He, S., Liu, J. D., & Zhang, G. (2026). Private firm information dissemination and analysts' public firm forecast accuracy. *Review of Accounting Studies, 31*, 1797–1836.

Clement, M. B. (1999). Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics, 27*(3), 285–303.

Clement, M. B., & Tse, S. Y. (2005). Financial analyst characteristics and herding behavior in forecasting. *The Journal of Finance, 60*(1), 307–341.

Davis, T., Lawson, J., & Nicholls, C. (2026). Does annual report complexity really impair a firm's information environment? *Review of Quantitative Finance and Accounting*.

Fu, S., Wang, K., Zhang, L., & Zheng, L. (2025). Bogging down investors: An unintended consequence of litigation risk. *Contemporary Accounting Research, 42*(2), 1045–1078.

Gow ID, Larcker DF, Reiss PC (2016) Causal inference in accounting research. *Journal of Accounting Research*, 54(2):477–523.

Ha, W., Han, S. Y., Lee, W. J., & Lim, Y. (2025). Discretionary reporting and analyst forecasts of operating income under IFRS. *The British Accounting Review*, 57(6), 101634.

Healy, P. M., & Palepu, K. G. (2001). Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature. *Journal of Accounting and Economics*, *31*(1-3), 405-440.

Kim, E. H., & Lyon, T. P. (2015). Greenwash vs. brownwash: Exaggeration and undue modesty in corporate sustainability disclosure. *Organization Science*, *26*(3), 705-723.

Kim, O., & Verrecchia, R. E. (1991). Trading volume and price reactions to public announcements. *Journal of Accounting Research, 29*(2), 302–321.

Kim, O., & Verrecchia, R. E. (1994). Market liquidity and volume around earnings announcements. *Journal of Accounting and Economics, 17*(1–2), 41–67.

Lang, M. H., & Lundholm, R. J. (1996). Corporate disclosure policy and analyst behavior. *The Accounting Review, 71*(4), 467–492.

Larcker DF, Rusticus TO (2010) On the use of instrumental variables in accounting research. *Journal of Accounting and Economics,* 49(3):186–205.

Lehavy, R., Li, F., & Merkley, K. (2011). The effect of annual report readability on analyst following and the properties of their earnings forecasts. *The Accounting Review*, *86*(3), 1087-1115.

Lev, B., & Sougiannis, T. (1996). The capitalization, amortization, and value-relevance of R&D. *Journal of Accounting and Economics*, *21*(1), 107-138.

Li, F. (2008). Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics, 45*(2–3), 221–247.

Liu, X. G., & Natarajan, R. (2012). The effect of financial analysts' strategic behavior on analysts' forecast dispersion. *The Accounting Review, 87*(6), 2123–2149.

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

Lu, Y., Harper, A., & Sarkar, S. K. (2026). How CEO stock option compensation shapes analyst coverage. *Review of Quantitative Finance and Accounting*, 66, 1539–1567.

McElheran K, Li JF, Brynjolfsson E, Kroff Z, Dinlersoz E, Foster L, Zolas N (2024) AI adoption in America: Who, what, and where. *Journal of Economics and Management Strategy,* 33(2):375–415

Mikhail, M. B., Walther, B. R., & Willis, R. H. (1997). Do security analysts improve their performance with experience? *Journal of Accounting Research, 35*, 131–157.

Pairolero, N. A., Giczy, A. V., Torres, G., Islam Erana, T., Finlayson, M. A., & Toole, A. A. (2025). The artificial intelligence patent dataset (AIPD) 2023 update. *The Journal of Technology Transfer, 50*, 2587–2610.

Park, M., & Zach, T. (2025). Analysts' forecasting models and uncertainty about the past. *Review of Accounting Studies*, 30, 2376–2418

Peters, R. H., & Taylor, L. A. (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics, 123*(2), 251–272.

Shanthikumar, D., & Yoo, I. S. (2026). Beyond automation: AI and the human value of sell-side analysts. *Journal of Accounting Research*.

Vo, H., & Pham, M. D. (2025). Beware of false prophets: Cybersecurity risk and strategic voluntary disclosure. *The British Accounting Review*, 101578.

Table 1. Descriptive Statistics and Univariate Comparisons by AI Washing

Panel A. Full Sample

Variable	N	Mean	SD	Min	Max
AIWashing	13,450	0.0000	0.4042	-0.9650	0.8820
FDISP	13,450	0.0136	0.0081	0.0004	0.0543
ForecastError	13,450	0.0261	0.0178	0.0010	0.1195
AINarrativeIntensity	13,450	0.0019	0.0017	0.0000	0.0115
ActualAIOutput	13,450	0.0221	0.0381	0.0000	0.4354
RDINT	13,450	0.0876	0.1104	0.0000	0.7952
SIZE	13,450	7.4548	1.4318	3.6000	11.8000
GROWTH	13,450	0.0852	0.1645	-0.4000	0.6500
SEGMENTS	13,450	1.1955	0.4973	0.0000	2.4849
INST	13,450	0.7145	0.1563	0.1219	0.9900
MFCOUNT	13,450	1.9992	1.4190	0.0000	10.0000
STDRET	13,450	0.0790	0.0351	0.0121	0.2800
10KNEWS	13,450	0.0383	0.0241	0.0033	0.2200
EARNVOL	13,450	0.0371	0.0216	0.0040	0.2200
LOSS	13,450	0.1624	0.3688	0.0000	1.0000
FOLLOW	13,450	2.3959	0.2607	1.0986	3.1781
FHORIZON	13,450	5.5686	0.1680	4.7095	5.8889
BOG	13,450	80.5773	11.5962	45.0000	121.3793

Panel B. Low versus High AIWashing

Variable	N Low	Low AIWashing	N High	High AIWashing	Difference (High−Low)
AIWashing	6,727	-0.3311	6,723	0.3314	0.6625***
FDISP	6,727	0.0139	6,723	0.0132	-0.0007***
ForecastError	6,727	0.0253	6,723	0.0269	0.0016***
AINarrativeIntensity	6,727	0.0011	6,723	0.0027	0.0016***
ActualAIOutput	6,727	0.0400	6,723	0.0042	-0.0358***
RDINT	6,727	0.0870	6,723	0.0882	0.0012
SIZE	6,727	7.5420	6,723	7.3675	-0.1745***
GROWTH	6,727	0.0844	6,723	0.0861	0.0017
SEGMENTS	6,727	1.1918	6,723	1.1993	0.0075
INST	6,727	0.7134	6,723	0.7156	0.0022
MFCOUNT	6,727	2.0036	6,723	1.9948	-0.0088
STDRET	6,727	0.0790	6,723	0.0791	0.0001
10KNEWS	6,727	0.0382	6,723	0.0384	0.0002
EARNVOL	6,727	0.0371	6,723	0.0371	-0.0001
LOSS	6,727	0.1655	6,723	0.1593	-0.0061
FOLLOW	6,727	2.3974	6,723	2.3944	-0.0030
FHORIZON	6,727	5.5664	6,723	5.5707	0.0043
BOG	6,727	80.6682	6,723	80.4863	-0.1819

Notes: Panel A reports descriptive statistics for the full sample. Panel B divides observations at the sample median of *AIWashing* and reports group means and their differences (High − Low).

AIWashing is the difference between the within-year percentile ranks of AI Narrative Intensity and subsequently verifiable Actual AI Output. Its mean is zero by construction. Continuous control variables are winsorized at the 1st and 99th percentiles. Variable definitions are provided in Appendix Table A2. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2. Baseline Results

Variables	FDISP	FDISP	FDISP
	(1)	(2)	(3)
AIWashing	-0.0011***	-0.0010***	-0.0011***
	(-4.201)	(-3.942)	(-4.052)
SIZE	-0.0003***	-0.0003***	-0.0003***
	(-6.624)	(-6.149)	(-6.196)
GROWTH	0.0015***	0.0013***	0.0013***
	(3.449)	(2.913)	(3.003)
RDINT	0.0017**	0.0015**	0.0014*
	(2.391)	(2.148)	(1.906)
SEGMENTS	0.0007***	0.0007***	0.0007***
	(4.910)	(4.633)	(4.577)
INST	-0.0019***	-0.0019***	-0.0020***
	(-4.526)	(-4.190)	(-4.299)
MFCOUNT	-0.0001**	-0.0001*	-0.0001*
	(-2.151)	(-1.748)	(-1.731)
STDRET	0.0146***	0.0162***	0.0158***
	(7.366)	(7.711)	(7.544)
10KNEWS	0.0054*	0.0055*	0.0060*
	(1.820)	(1.721)	(1.856)
EARNVOL	0.0186***	0.0186***	0.0185***
	(5.623)	(5.224)	(5.236)
LOSS	0.0004**	0.0006***	0.0006***
	(2.137)	(3.038)	(2.968)
FOLLOW	-0.0002	-0.0001	-0.0001
	(-0.717)	(-0.431)	(-0.315)
FHORIZON	0.0013***	0.0013***	0.0013***
	(3.202)	(3.102)	(3.096)
BOG	0.0001***	0.0001***	0.0001***
	(10.296)	(9.303)	(9.337)
Constant	0.0034***	0.0047***	0.0029**
	(3.080)	(3.891)	(2.349)
Firm FE	No	Yes	Yes
Year FE	Yes	Yes	Absorbed
Industry × Year FE	No	No	Yes
N	13,450	13,450	13,450
Adj. R^2	0.0266	0.0342	0.0355

Notes: This table reports OLS regressions of *FDISP* on *AIWashing*. Column (1) includes year fixed effects, Column (2) additionally includes firm fixed effects, and Column (3) includes firm and industry-by-year fixed effects. All specifications include the baseline controls defined in Appendix Table A2. Firm-clustered t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 3. Analyst Forecast Convergence around the 10-K Filing

Variables	Full window	±30 days	Stable analyst set
	(1)	(2)	(3)
AIWashing × Post	-0.0002***	-0.0002**	-0.0002***
	(-2.786)	(-2.174)	(-3.061)
Controls	Absorbed	Absorbed	Absorbed
Firm-Year FE	Yes	Yes	Yes
Event-Window FE	Yes	Yes	Yes
N	75,436	51,178	65,298
Within R^2	0.0824	0.0741	0.0879

Notes: This table examines changes in analyst forecast dispersion around the focal 10-K filing. Column (1) uses six event windows: [−60,−31], [−30,−11], [−10,−1], [0,+10], [+11,+30], and [+31,+60]. Column (2) restricts the analysis to the ±30-trading-day region, and Column (3) requires a stable analyst set across the relevant pre- and post-filing snapshots. A firm-year-window is retained only when at least two distinct analysts have valid outstanding FY1 forecasts. All specifications include firm-year and event-window fixed effects. Firm-clustered t-statistics are reported in parentheses.

Table 4. Instrumental-Variable Analysis (2SLS)

Variables	First Stage: AIWashing	Second Stage: FDISP
	(1)	(2)
LowAIExposure	0.0794***	
	(6.551)	
AIWashing (Instrumented)		-0.0052**
		(-2.067)
TopUniExposure	-0.0103	-0.0011
	(-0.364)	(-0.421)
CSStrongExposure	-0.0765***	-0.0023*
	(-2.831)	(-1.782)
Constant	0.0918***	0.0151***
	(3.912)	(4.338)
Controls	Yes	Yes
Industry FE	Yes	Yes
Year FE	Yes	Yes
N	10,286	10,286
Centered R^2	0.0192	0.0217
Partial R^2	0.0121	—
Kleibergen-Paap rk Wald F	42.92	—

Notes: This table reports two-stage least squares (2SLS) estimates using *LowAIExposure*, constructed from firms' pre-sample exposure to AI-strong universities through historical talent networks, as the excluded instrument for *AIWashing*. The IV sample is restricted to firms for which the pre-sample STEM employee–university network and the corresponding university research measures can be reliably constructed. Firms without observable historical network information are excluded rather than assigned zero exposure. Column (1) reports the first-stage estimates, and Column (2) reports the second-stage estimates with *FDISP* as the dependent variable. Both stages include *TopUniExposure*, *CSStrongExposure*, the full baseline controls (including RDINT), industry fixed effects, and year fixed effects. Firm fixed effects are excluded because the instrument is time-invariant. The partial R^2 and Kleibergen–Paap rk Wald F-statistic assess instrument relevance. Firm-clustered t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 5. AI Washing and Forecast Accuracy

Variables	Consensus Forecast Error	Analyst Forecast Error
	(1)	(2)
AIWashing	0.0020***	0.0018***
	(3.151)	(3.213)
Constant	0.0108***	0.0126***
	(4.015)	(4.288)
Controls	Yes	Yes
Firm FE	Yes	Yes
Analyst FE	No	Yes
Year FE	Yes	Yes
N	13,450	124,751
Adj. R^2	0.0813	0.0672

Notes: This table reports regressions of forecast errors on *AIWashing*. Column (1) examines absolute consensus forecast error, while Column (2) examines individual analyst forecast error. Both specifications include the full baseline controls, firm fixed effects, and year fixed effects. Column (2) additionally includes analyst fixed effects. Standard errors are clustered at the firm level in Column (1) and two-way clustered by firm and analyst in Column (2). Corresponding t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 6. Subsequent Forecast Correction

Variables	Revision Magnitude	Correction toward Realized EPS
	(1)	(2)
AIWashing	0.0006***	0.0087**
	(2.884)	(2.243)
Constant	0.0041***	0.4726***
	(3.025)	(31.422)
Controls	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
N	11,872	11,872
Adj. R^2	0.0398	0.0615

Notes: This table examines subsequent revisions to analysts' initial post-10-K FY1 consensus forecasts. *Revision Magnitude* is the absolute difference between the initial post-filing consensus and the last distinct valid consensus preceding the corresponding earnings announcement for the same target fiscal period, scaled by the same price denominator used for *FDISP*. Correction toward *Realized EPS* equals one if the later consensus is closer to realized EPS than the initial consensus, and zero otherwise. Firm-years without both valid consensus snapshots are excluded. Both specifications include the full baseline controls, firm fixed effects, and year fixed effects.

Table 7. Analyst-Level Convergence toward the Prevailing Consensus

Variables	Post-Filing Distance	Change in Distance
	(1)	(2)
AIWashing	-0.0002***	-0.0002***
	(-3.313)	(-3.052)
Pre-Filing Distance	0.7114***	
	(3.182)	
Constant	0.0018***	-0.0003
	(5.227)	(-1.124)
Controls	Yes	Yes
Firm FE	Yes	Yes
Analyst FE	Yes	Yes
Year FE	Yes	Yes
N	124,751	124,751
Adj. R^2	0.4217	0.0935

Notes: This table examines individual analysts' convergence toward the prevailing consensus. *Distance to Consensus* is defined as the absolute difference between an analyst's outstanding FY1 EPS forecast and the contemporaneous leave-one-out consensus (the mean forecast of all other analysts covering the same firm), scaled by the same price denominator used for *FDISP*. *Pre-Filing Distance* is measured using the last valid consensus before the focal 10-K filing, whereas *Post-Filing Distance* uses the first valid post-filing consensus. Change in Distance equals *Post-Filing Distance* minus *Pre-Filing Distance*, with negative values indicating convergence. Analysts must have valid forecasts in both periods. Column (1) controls for Pre-Filing Distance, while Column (2) uses Change in Distance as the dependent variable. Both specifications include the full baseline controls, firm, analyst, and year fixed effects. Two-way clustered t-statistics at the firm and analyst levels are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 8. Analyst Herding and Forecast Boldness

Variables	HERD	BOLD
	(1)	(2)
AIWashing	0.0085**	-0.0066*
	(2.312)	(-1.802)
Constant	0.4217***	0.2861***
	(24.188)	(17.445)
Controls	Yes	Yes
Firm FE	Yes	Yes
Analyst FE	Yes	Yes
Year FE	Yes	Yes
N	102,384	102,384
Adj. R^2	0.1128	0.0974

Notes: This table examines analyst herding and forecast boldness following Clement and Tse (2005). *HERD* and *BOLD* are defined using each analyst's first actual post-filing FY1 forecast revision relative to the analyst's preceding forecast and the prevailing pre-revision leave-one-out consensus. Revisions equal to either benchmark are retained and coded as zero for both indicators; consequently, *HERD* and *BOLD* are not complementary binary variables. Both specifications include the full baseline controls, firm, analyst, and year fixed effects. Standard errors are two-way clustered by firm and analyst, with corresponding t-statistics in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 9. AI Washing and Analysts' Common versus Private Information

Variables	Consensus (ρ)	Common Information Precision (h)	Private Information Precision (s)
	(1)	(2)	(3)
AIWashing	0.0612***	0.0187	-0.0478**
	(3.004)	(0.842)	(-2.271)
Constant	0.0043	0.0019	-0.0069
	(0.181)	(0.076)	(-0.277)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	12,987	12,987	12,987
Adj. R^2	0.1182	0.1091	0.1047

Notes: This table examines analysts' information environments following Barron et al. (1998). *D* denotes the cross-sectional variance of price-scaled individual forecasts, *SE* the squared price-scaled consensus forecast error, and *N* the number of analysts. Common uncertainty is $C = SE - D/N$, and total uncertainty is $V = C + D$. Columns (1)–(3) examine consensus ($\rho = C/V$), common information precision ($h = C/V^2$), and private information precision ($s = D/V^2$), respectively, under the equal-private-information-precision assumption. The analysis retains only observations with strictly positive model-implied common uncertainty C and total uncertainty V, and for which all three information measures are finite and defined. The three dependent variables are standardized within this valid analysis sample before estimation. All specifications include the full baseline controls, firm fixed effects, and year fixed effects. Firm-clustered t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 10. Firm-Specific Analyst Experience and Susceptibility to AI Washing

Variables	Change in Distance	HERD
	(1)	(2)
AIWashing	-0.0002***	0.0101**
	(-3.422)	(2.424)
High Firm Experience	-0.0000	-0.0021
	(-0.684)	(-0.812)
AIWashing × High Firm Experience	0.0001**	-0.0058**
	(2.186)	(-2.034)
Constant	-0.0003	0.4231***
	(-1.004)	(23.962)
Controls	Yes	Yes
Firm FE	Yes	Yes
Analyst FE	Yes	Yes
Year FE	Yes	Yes
N	124,751	102,384
Adj. R^2	0.0951	0.1136

Notes: This table examines the moderating role of analysts' firm-specific coverage experience. *High Firm Experience* equals one if an analyst's coverage tenure for the focal firm, measured before the 10-K filing, exceeds the sample-year median, and zero otherwise. Column (1) uses *Change in Distance*, defined as *Post-Filing Distanc*e minus *Pre-Filing Distance*, while Column (2) uses *HERD*, an indicator of forecast revisions toward the prevailing leave-one-out consensus. Both specifications include *AIWashing*, *High Firm Experience*, their interaction, the full baseline controls, and firm, analyst, and year fixed effects. Two-way clustered t-statistics at the firm and analyst levels are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 11. AI Narrative–Patent-Output Alignment and Analyst Forecast Outcomes

Variables	FDISP	Forecast Error	Revision Magnitude
	(1)	(2)	(3)
AI Washing	-0.0008***	0.0019***	0.0006**
	(-3.316)	(3.617)	(2.122)
Substantive AI	-0.0006***	-0.0015***	-0.0004*
	(-2.861)	(-3.892)	(-1.807)
Unexpected AI Output	0.0001	-0.0007*	-0.0003
	(0.531)	(-1.839)	(-1.403)
Constant	0.0046***	0.0109***	0.0040***
	(3.617)	(3.948)	(3.071)
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	13,450	13,450	11,872
Adj. R^2	0.0352	0.0831	0.0416

Notes: This table distinguishes four mutually exclusive groups using a within-year median split of AI Narrative Intensity and an indicator for positive AI-related pre-grant patent output. High Narrative equals one when AI Narrative Intensity is strictly above the within-year median; observations equal to the median are classified as Low Narrative. Positive AI Output equals one when at least one AI-classified pre-grant patent publication is observed during the verification window. *Low Narrative–Zero AI Patent Output* (low narrative, zero AI output) is the omitted reference group. *High Narrative–Zero AI Patent Output* denotes high narrative and zero AI output; *High Narrative–Positive AI Patent Output* denotes high narrative and positive AI output; and *Low Narrative–Positive AI Patent Output* denotes low narrative and positive AI output. The categorical *AI Washing* group is distinct from the continuous AIWashing rank-gap measure used in the baseline analyses. Columns (1)–(3) examine FDISP, Forecast Error, and Revision Magnitude, respectively. All specifications include the full baseline controls, firm fixed effects, and year fixed effects. Firm-clustered t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 12. AI Washing and Earnings-Announcement Reactions

Variables	|CAR[-1,+1]|	|CAR[-1,+1]|
	(1)	(2)
AIWashing	0.0014**	0.0012**
	(2.312)	(2.064)
ForecastError		0.1132***
		(4.751)
Constant	0.0191***	0.0162***
	(5.107)	(4.388)
Controls	Yes	Yes
Firm FE	Yes	Yes
Year FE	Yes	Yes
N	13,450	13,450
Adj. R^2	0.0412	0.0527

Notes: This table examines the association between *AIWashing* and subsequent earnings-announcement reactions. The dependent variable, $|CAR[-1,+1]|$, is the absolute cumulative market-adjusted return over the [−1,+1] trading-day window surrounding the subsequent earnings announcement. Column (1) reports the baseline association, while Column (2) additionally includes *ForecastError*, defined as the absolute difference between the initial post-10-K consensus FY1 EPS forecast and realized EPS, scaled by the same price denominator used for *FDISP*. Column (2) represents a conditional specification rather than a causal mediation test. Both specifications include the full baseline controls, firm fixed effects, and year fixed effects. Firm-clustered t-statistics are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Figure 1. AI Narrative Intensity and Subsequently Verifiable AI Output

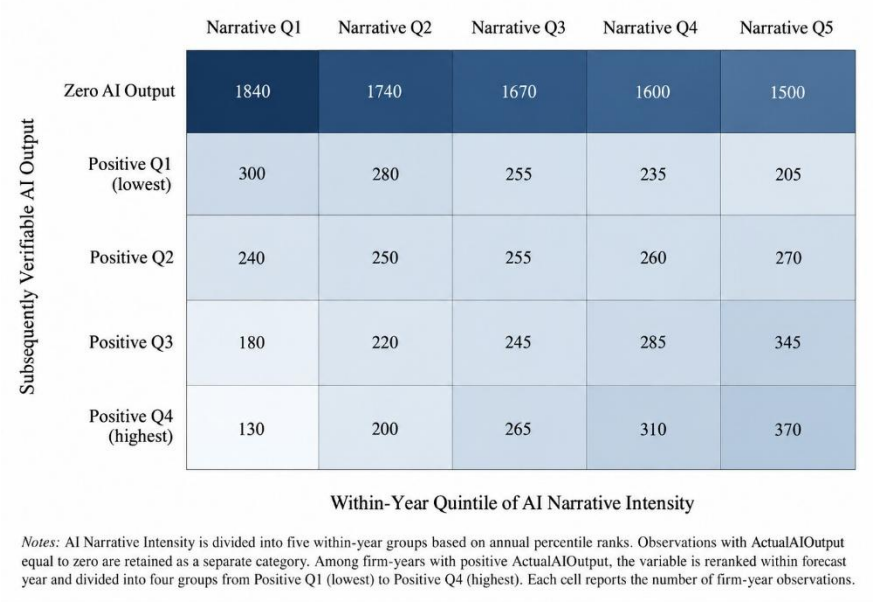

Notes: This figure presents the joint distribution of AI Narrative Intensity and subsequently verifiable AI Output. AI Narrative Intensity is divided into five within-year groups based on annual percentile ranks. Observations with *ActualAIOutput* equal to zero are retained as a separate category. Among firm-years with positive *ActualAIOutput*, the variable is reranked within forecast year and divided into four groups, from Positive Q1 (lowest) to Positive Q4 (highest). Each cell reports the number of firm-year observations.

Figure 2. Analyst Forecast Dispersion around the 10-K Filing

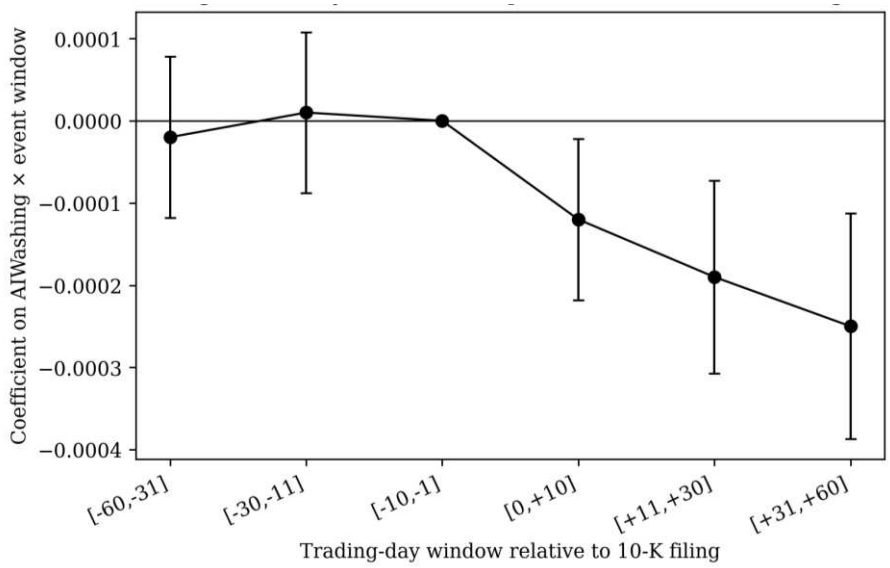

Notes: This figure plots the coefficients on interactions between *AIWashing* and event-window indicators around the focal 10-K filing. The [−10, −1] trading-day window serves as the omitted reference period. Vertical bars represent 95% confidence intervals based on firm-clustered standard errors.

Figure 3. AI Washing, Forecast Disagreement, and Forecast Error

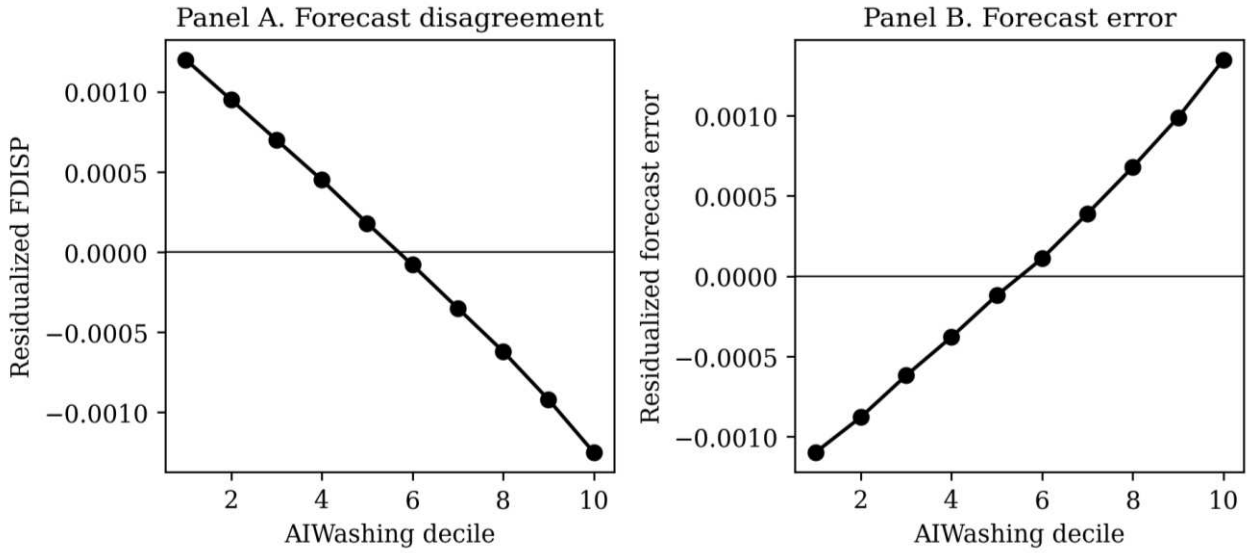

Notes: This figure plots the mean residualized forecast outcomes for each AIWashing decile. Panel A reports analyst forecast disagreement (*FDISP*), and Panel B reports absolute consensus forecast error (*ForecastError*). Both outcomes are residualized using the full baseline controls, firm fixed effects, and year fixed effects. Points represent within-decile averages of the residualized outcomes.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryAppendix.docx